

Institutional Expanded Evaluation Brief

Company: **Databricks**

Sector: **AI Infrastructure / Enterprise Data Platform**

Region: **Global**

Prepared by: **Cohres Ltd**

Date: **July 2026**

Institutional Expanded Evaluation Brief

Expanded Evaluation

Prepared using publicly available information and Cohres analysis. This evaluation is intended to validate key commercial assumptions and determine whether full commercial due diligence is warranted.

TABLE OF CONTENTS

Executive Summary	Page 2
Company Overview	Page 5
Team Assessment	Page 7
Market Analysis	Page 9
Product & Platform Positioning	Page 12
Competitive Landscape	Page 15
Business Model	Page 18
Key Strengths, Risks & Diligence Priorities	Page 22
Transaction Landscape & IPO Considerations	Page 27
Conviction Assessment	Page 31
Sources & Disclaimer	Page 33



EXECUTIVE SUMMARY

Company Overview

[Databricks](#) is an enterprise AI and data platform built around the Lakehouse architecture, enabling organizations to unify data engineering, analytics, machine learning, governance, and artificial intelligence within a single operating environment. Rather than competing as a standalone analytics or database provider, the company positions itself as foundational infrastructure supporting enterprise data management and AI deployment across large organizations.

Since its founding in 2013 by the creators of Apache Spark, Databricks has expanded from an open-source data processing company into one of the largest private enterprise software businesses globally. Public disclosures indicate the company serves thousands of enterprise customers across financial services, healthcare, manufacturing, retail, technology, and the public sector. Recent funding rounds valuing the company at more than \$60B reflect investor conviction that enterprise AI adoption will increasingly depend on scalable data infrastructure rather than foundation models alone.

The platform combines data warehousing, data engineering, streaming analytics, machine learning, governance, and AI development into a unified environment designed to reduce data fragmentation and accelerate enterprise AI deployment. Strategic partnerships with Microsoft, AWS, Google Cloud, NVIDIA, Anthropic, and OpenAI further strengthen Databricks' position within the broader AI infrastructure ecosystem.

Core Thesis

The investment case depends on whether Databricks becomes mission-critical enterprise infrastructure underpinning data, analytics, governance, and AI rather than remaining primarily a high-growth software platform.

The remainder of this report evaluates whether Databricks' market position, platform architecture, business model, competitive durability, and public-market readiness support progression to full Commercial Due Diligence.

While generative AI has accelerated enterprise investment across the technology stack, sustainable competitive advantage is increasingly shifting toward organizations capable of managing, governing, and operationalizing proprietary enterprise data at scale. Foundation models may become increasingly commoditized, but enterprise data infrastructure, governance, and production deployment remain significantly more difficult to replicate.

Databricks is well positioned to benefit from this structural transition. The company's unified Lakehouse architecture, expanding AI capabilities, broad ecosystem partnerships, and deep enterprise relationships position it as an infrastructure provider enabling organizations to build, deploy, and manage AI applications using their own proprietary data.

The strongest strategic advantage does not appear to be the underlying AI models themselves. Instead, differentiation increasingly comes from becoming the platform where enterprise data is stored, governed, processed, and transformed into production AI systems. If enterprise customers continue consolidating data, analytics, and AI workflows onto Databricks, the platform may develop durable switching costs through operational integration, data governance, and ecosystem dependency.

At the same time, long-term defensibility remains dependent on Databricks continuing to differentiate beyond its open-source foundation while successfully defending its position against increasingly integrated hyperscaler platforms. Cloud hyperscalers, Snowflake, Microsoft, Oracle, Google Cloud, Amazon Web Services, and emerging AI infrastructure providers continue expanding their own data platform capabilities. The long-term investment outcome will depend on whether Databricks can maintain technological leadership while expanding

platform adoption faster than competing ecosystems converge around similar enterprise AI functionality.

Key Strengths

- Market-leading position within enterprise AI and data infrastructure
- Large and expanding enterprise customer base across multiple industries
- Unified Lakehouse architecture integrating data engineering, analytics, governance, and AI
- Strong ecosystem partnerships with major cloud providers and frontier AI companies
- Significant developer adoption supported by Apache Spark and open-source credibility
- Rapid revenue growth supported by increasing enterprise AI investment
- Strong institutional backing and financial resources supporting continued expansion
- Attractive positioning ahead of a potential public market listing

Key Risks

- Intensifying competition from Snowflake, hyperscalers, and integrated cloud platforms
- Continued pricing pressure as enterprise AI infrastructure becomes increasingly competitive
- Execution risk associated with maintaining innovation leadership across multiple product categories
- Dependence on sustained enterprise AI spending and cloud infrastructure investment
- High valuation requiring continued rapid revenue growth to justify investor expectations
- Increasing platform overlap between data infrastructure, AI platforms, and cloud-native services
- Potential commoditization of certain analytics and AI capabilities over time
- Complexity of enterprise implementations and long sales cycles
- Enterprise implementation complexity may slow customer adoption and delay time-to-value, particularly outside highly sophisticated data organizations.

Core Investment Questions

1. Can Databricks become the dominant operating platform for enterprise AI and data infrastructure across large organizations?
2. Are enterprise customers consolidating workloads onto Databricks in ways that create durable switching costs and long-term retention?
3. Can the company sustain differentiation as cloud providers and competitors continue integrating similar AI and data platform capabilities?
4. Will expanding AI workloads accelerate platform adoption and revenue growth beyond traditional analytics use cases?
5. Does the current valuation appropriately reflect the company's long-term competitive position and potential public market trajectory?

Conviction Assessment

Relative to other enterprise AI infrastructure companies, Databricks appears to possess one of the strongest combinations of technological leadership, enterprise adoption, ecosystem partnerships, and commercial momentum. Public evidence suggests the company is increasingly evolving beyond a data analytics platform into foundational infrastructure supporting enterprise AI development, governance, and deployment.

The principal investment question is no longer whether organizations require AI infrastructure, but whether Databricks can establish sufficient platform centrality to become an indispensable component of enterprise technology architecture. If successful, the company could benefit from durable customer relationships, expanding workloads, and increasing operational dependency as enterprises scale AI initiatives.

However, long-term success remains dependent on continued execution, competitive differentiation, and the ability to maintain platform leadership within an increasingly crowded enterprise AI infrastructure market.



Further diligence should therefore focus on customer retention, workload expansion, pricing durability, competitive positioning, infrastructure economics, and the sustainability of the company's long-term growth trajectory.

Based on currently available information, Databricks warrants progression beyond Expanded Evaluation and into full Commercial Due Diligence. This assessment should focus on validating revenue quality, customer expansion dynamics, competitive durability, unit economics, and long-term platform defensibility before an investment recommendation is made.

Methodology

This brief represents a commercial evaluation based on publicly available information, company materials, investor commentary, customer case studies, industry analysis, and Cohres assessment. Several conclusions represent commercial interpretation and workflow analysis rather than verified internal operating metrics.

Public information does not provide full visibility into retention metrics, expansion economics, implementation costs, gross margins, or long-term cohort durability.



COMPANY OVERVIEW

Company Background

Databricks is an enterprise AI and data platform that enables organizations to unify data engineering, analytics, machine learning, governance, and artificial intelligence within a single cloud-native environment. Founded in 2013 by the creators of Apache Spark at the University of California, Berkeley, the company was established to commercialize large-scale distributed data processing while simplifying how enterprises manage and derive value from data.

Over the past decade, Databricks has expanded well beyond its origins as a big data analytics platform. The company now positions itself as a comprehensive enterprise AI platform, enabling organizations to build, govern, and deploy AI applications using proprietary enterprise data. Its strategic focus has shifted from improving analytical workflows to becoming foundational infrastructure supporting enterprise-wide digital transformation and artificial intelligence adoption.

Platform Evolution

Databricks has progressively expanded its platform through a series of strategic product developments designed to consolidate fragmented enterprise data environments.

The company initially gained market adoption through Apache Spark, which became one of the most widely adopted distributed data processing frameworks globally. Building upon this foundation, Databricks introduced the Lakehouse architecture, combining the scalability and flexibility of data lakes with the governance, reliability, and performance traditionally associated with enterprise data warehouses.

More recently, the platform has evolved into what Databricks describes as a Data Intelligence Platform, integrating data engineering, SQL (Structured query language) analytics, streaming, machine learning, governance, data sharing, and generative AI capabilities into a unified operating environment.

This progression reflects a broader strategic objective of reducing complexity across enterprise data infrastructure while providing organizations with a single platform capable of supporting both traditional analytics and emerging AI workloads.

Commercial Scale & Enterprise Adoption

Databricks has established itself as one of the largest privately held enterprise software companies globally. Public disclosures indicate the platform serves thousands of enterprise customers, including a significant proportion of Fortune 500 organizations, across multiple regulated and data-intensive industries.

Commercial adoption has accelerated alongside growing enterprise investment in cloud computing, advanced analytics, and artificial intelligence. Organizations increasingly utilize Databricks to consolidate data management, support large-scale analytics, develop machine learning models, and operationalize generative AI applications using proprietary business data.

Technology & Product Ecosystem

Databricks has developed an integrated technology ecosystem designed to support the complete enterprise data and AI lifecycle.

The platform is supported by several strategic technologies, including Delta Lake, Unity Catalog, MLflow, Databricks SQL, and Mosaic AI. Together, these components enable organizations to manage structured and unstructured data, maintain governance controls, develop AI models, and deploy production-scale AI applications from within a unified environment.



Databricks also maintains strategic partnerships with Microsoft Azure, Amazon Web Services, Google Cloud, NVIDIA, Anthropic, and OpenAI. These relationships strengthen the platform's position within the broader AI infrastructure ecosystem while allowing customers to deploy AI workloads across multiple cloud environments using leading foundation models.

Rather than competing directly as a foundation model developer, Databricks increasingly positions itself as the infrastructure layer that enables enterprises to operationalize AI securely, govern proprietary data, and scale production AI systems across the organization.

Capitalization & Strategic Positioning

Databricks has completed multiple institutional funding rounds, attracting investment from leading venture capital firms, growth equity investors, strategic corporate investors, and sovereign wealth funds. Recent financing rounds have valued the company at more than \$60B, making it one of the highest-valued private software companies globally.

Major investors include Thrive Capital, Andreessen Horowitz, DST Global, GIC, Insight Partners, WCM Investment Management, Ontario Teachers', ICONIQ Growth, MGX, Sands Capital, and Wellington Management, reflecting broad support from venture capital, institutional investors, strategic partners, and sovereign wealth funds.

The investment significance is not simply the amount of capital raised, but what that capital is underwriting: the possibility that Databricks becomes foundational enterprise AI infrastructure through which enterprises organize, govern, and operationalize data for AI. Capital has been deployed to accelerate product development, expand internationally, strengthen enterprise sales capabilities, pursue strategic acquisitions, and deepen investments in artificial intelligence infrastructure.

The company also benefits from a strong ecosystem of technology partners, enterprise customers, and open-source communities that reinforce its market position. Strategic collaborations with cloud hyperscalers and leading AI providers increase platform adoption while extending Databricks' role within enterprise technology stacks.

The scale of commercial growth, institutional backing, ecosystem partnerships, and product expansion suggests that Databricks is being positioned not simply as a provider of data analytics software, but as a foundational enterprise AI infrastructure platform supporting the next generation of data-driven and AI-enabled organizations.

TEAM ASSESSMENT

Management Overview

Databricks was founded in 2013 by Ali Ghodsi, Reynold Xin, Patrick Wendell, Matei Zaharia, Ion Stoica, Scott Shenker, Arsalan Tavakoli-Shiraji, and Andy Konwinski, many of whom were responsible for creating Apache Spark at the University of California, Berkeley. The founding team combined deep expertise in distributed computing, data systems, and large-scale infrastructure with extensive academic research in cloud computing and machine learning.

Since its founding, Databricks has evolved from a research-driven startup into one of the world's largest private enterprise software companies. As the organization has expanded globally, the leadership team has broadened beyond technical founders to include experienced executives across enterprise sales, finance, product management, customer success, and international operations.

The key management question is no longer whether the team can build innovative technology. It is whether the organization can continue executing at global scale while maintaining product leadership across an increasingly competitive enterprise AI infrastructure market.

Founder - Market Fit

Few enterprise software companies demonstrate a stronger founder-market fit than Databricks.

The company's founders were directly responsible for developing Apache Spark, one of the most widely adopted distributed computing frameworks globally. Their technical expertise enabled Databricks to commercialize technologies that addressed increasingly complex enterprise data processing challenges while maintaining strong credibility within both academic and commercial communities.

Rather than entering an established software category, the founders helped define many of the underlying technologies that now support modern enterprise analytics and AI infrastructure. This provides Databricks with an unusually strong technical foundation and long-term product vision.

Technical & Product Execution Capability

Databricks has consistently demonstrated the ability to expand its platform beyond its original data processing capabilities while maintaining technological leadership across multiple product categories.

The company's product evolution from Apache Spark to the Lakehouse architecture and, more recently, the Data Intelligence Platform reflects an ability to anticipate changing enterprise technology requirements. Product expansion into governance, machine learning, data sharing, generative AI, and AI application development suggests disciplined execution rather than incremental feature development.

One of Databricks' greatest competitive strengths is its ability to integrate rapidly evolving AI technologies into an existing enterprise platform rather than building isolated AI products.

Enterprise Scaling & Organizational Development

Databricks has successfully transitioned from an engineering-led startup into a global enterprise software organization serving many of the world's largest companies.

Its leadership team has also demonstrated an ability to attract experienced executives with backgrounds in enterprise software, cloud infrastructure, and global business operations. This organizational development has enabled Databricks to scale commercial operations while continuing significant investment in research and product innovation.



As the company approaches potential public market readiness, continued organizational discipline, financial management, and operational execution will become increasingly important.

Customer Enablement & Implementation Capability

Unlike many enterprise software platforms that focus primarily on technology deployment, Databricks has invested heavily in customer enablement, training, consulting partnerships, and technical support to facilitate enterprise adoption.

Long-term commercial success will depend not only on technological innovation but also on the organization's ability to consistently deliver successful enterprise transformations across diverse industries.

Key Team Strengths

Key team strengths include:

- founders with exceptional technical credibility and deep expertise in distributed computing and AI infrastructure
- proven ability to commercialize open-source innovation into enterprise software platforms
- strong product execution demonstrated through continuous platform expansion
- experienced executive leadership across enterprise software, cloud infrastructure, and global operations
- ability to attract leading engineering talent and strategic technology partners
- strong relationships with major cloud providers and enterprise customers
- demonstrated capability to scale globally while maintaining rapid product innovation

Key Risks

Key team risks include:

- maintaining innovation leadership as the organization continues to grow in size and complexity
- balancing long-term platform development with increasing public market expectations following a potential IPO
- managing competition for technical talent within the rapidly expanding AI sector
- sustaining execution velocity while coordinating an increasingly broad product portfolio
- successfully integrating acquisitions and emerging AI technologies into the broader platform
- maintaining organizational alignment across global operations during continued international expansion

Team Assessment

Databricks has one of the strongest technical leadership profiles within enterprise AI infrastructure, supported by founder-market fit, open-source credibility, and demonstrated product execution.

The key diligence question is no longer whether the team can build category-defining technology. It is whether management can sustain commercial execution as Databricks transitions from founder-led product innovation toward global enterprise sales, public company discipline, and increasingly competitive AI infrastructure markets.

Commercial Due Diligence should therefore focus on enterprise sales execution, organizational scalability, product prioritization, and the company's ability to maintain innovation velocity while operating at public-market scale.

MARKET ANALYSIS

Industry Context

Enterprise artificial intelligence is fundamentally reshaping how organizations manage, analyze, and utilize data. While early cloud adoption focused primarily on infrastructure modernization and cost efficiency, the rapid advancement of machine learning and generative AI has shifted strategic priorities toward extracting greater value from proprietary enterprise data.

As organizations accelerate AI adoption, the limiting factor is increasingly becoming data readiness rather than model availability. Foundation models have become more accessible through cloud providers and AI vendors, but enterprise success depends on the ability to integrate, govern, secure, and operationalize proprietary datasets across increasingly complex technology environments.

This structural shift has elevated enterprise data platforms from supporting infrastructure to strategic business assets. Organizations are increasingly seeking unified environments capable of managing data engineering, analytics, governance, and AI development within a single platform rather than maintaining fragmented technology stacks.

Market Size & Growth

Databricks operates across several of the fastest-growing segments of enterprise software, including cloud data platforms, AI infrastructure, analytics, machine learning, and enterprise data management.

Quantitatively, the market backdrop remains supportive. IDC estimates worldwide AI spending will increase from approximately \$235B to more than \$631B by 2028, implying significant continued investment across AI-enabled applications, infrastructure, and services. Gartner also forecasts worldwide public cloud end-user spending of \$723.4B in 2025, up from \$595.7B in 2024, reinforcing the scale of the infrastructure layer supporting enterprise AI adoption.

Global enterprise investment in artificial intelligence continues to accelerate as organizations prioritize digital transformation, automation, predictive analytics, and AI-enabled decision-making. Cloud migration remains a significant long-term driver, with enterprises continuing to modernize legacy infrastructure while increasing spending on scalable cloud-native data platforms.

In parallel, demand for generative AI applications has expanded the addressable market beyond traditional analytics. Organizations increasingly require platforms capable of supporting AI model development, retrieval-augmented generation (RAG), agentic AI workflows, real-time data processing, governance, and enterprise-scale deployment.

These trends suggest that enterprise AI infrastructure spending is likely to remain a structural growth market over the coming decade.

The key underwriting question is not whether the market grows, but whether enterprise AI spending consolidates around a small number of strategic infrastructure platforms or remains fragmented across hyperscalers, data platforms, internal tools, and specialist AI vendors.

Industry Structure

The enterprise AI infrastructure market consists of several overlapping technology segments that are progressively converging.

The market is increasingly shifting toward unified platforms capable of consolidating these functions into integrated operating environments.

This convergence creates opportunities for providers capable of reducing architectural complexity while supporting a broad range of enterprise data and AI workloads through a single platform.

However, industry convergence also increases competitive intensity as software vendors expand beyond their historical product categories.

Key Growth Drivers

Several structural trends are supporting continued expansion of enterprise AI infrastructure.

- **Enterprise AI Adoption:** Organizations are rapidly moving beyond experimentation toward deploying AI across core business processes, increasing demand for production-grade infrastructure capable of supporting large-scale AI applications.
- **Cloud Modernization:** Continued migration from on-premise infrastructure to public cloud environments is expanding demand for cloud-native data management and analytics platforms.
- **Data Growth:** The volume of enterprise data continues to grow across structured, unstructured, streaming, and multimodal formats, increasing demand for scalable storage, processing, and governance capabilities.
- **Regulatory Requirements:** Growing regulatory expectations surrounding data governance, privacy, security, and responsible AI are increasing the importance of enterprise platforms capable of managing compliance across complex environments.
- **Platform Consolidation:** Organizations increasingly seek to simplify technology stacks by reducing the number of vendors managing data, analytics, governance, and AI development, creating opportunities for integrated platforms.

Competitive Dynamics

Competition within enterprise AI infrastructure remains intense and continues to evolve rapidly.

Cloud hyperscalers, enterprise software vendors, database providers, analytics companies, and AI-native platforms are all expanding their capabilities in an effort to capture a larger share of enterprise AI spending.

Large cloud providers benefit from extensive infrastructure ownership, existing enterprise relationships, and integrated cloud ecosystems. Traditional database and analytics vendors continue enhancing their AI capabilities, while newer AI-native companies seek to differentiate through innovation, developer adoption, and platform flexibility.

As the market matures, competitive differentiation is expected to rely less on individual AI models and more on platform integration, governance, ecosystem breadth, customer adoption, and the ability to operationalize AI at enterprise scale.

Market Attractiveness

Several characteristics make enterprise AI infrastructure an attractive long-term market.

The market benefits from strong structural growth, increasing enterprise adoption, high switching costs, recurring software revenue models, and expanding AI investment across virtually every major industry.

Enterprise customers typically make long-term infrastructure decisions involving complex implementation processes, extensive data migration, governance requirements, and integration across existing technology environments. Once deployed successfully, these platforms often become deeply embedded within organizational operations, creating durable customer relationships and opportunities for expansion over time.

At the same time, the market is characterized by rapid technological change and significant competitive pressure. Vendors must continuously invest in product innovation to remain competitive while responding to evolving customer requirements and advances in AI technologies.

Market Assessment

The enterprise AI infrastructure market is attractive due to structural growth in AI adoption, cloud modernization, data governance, and platform consolidation.

For Databricks, the key market question is not whether demand grows, but whether enterprise AI spending consolidates around a small number of strategic infrastructure platforms. The strength of the investment thesis depends on Databricks capturing a disproportionate share of this consolidation rather than becoming one of many interchangeable data and AI platforms.

PRODUCT & PLATFORM POSITIONING

Product Overview

Databricks positions itself as a unified enterprise platform designed to manage the complete lifecycle of data, analytics, machine learning, and artificial intelligence. Rather than offering individual point solutions for data engineering, business intelligence, or AI development, the platform integrates these capabilities into a single cloud-native environment intended to reduce fragmentation across enterprise technology stacks.

The company's strategic objective extends beyond improving analytical productivity. Databricks increasingly seeks to become the operational platform through which organizations ingest, govern, analyze, and operationalize proprietary data to support production AI systems. This positioning places the platform at the center of enterprise digital transformation initiatives rather than at the edge of existing data infrastructure.

Unified Data & AI Platform

One of Databricks' most significant competitive advantages is its ability to consolidate historically separate enterprise technologies into a unified operating environment.

Traditional enterprise architectures frequently require organizations to manage multiple vendors across data storage, analytics, governance, machine learning, and business intelligence. These fragmented environments increase implementation complexity, duplicate data movement, and create operational inefficiencies.

Databricks addresses this challenge through its Lakehouse architecture, enabling organizations to manage structured and unstructured data while supporting analytical workloads and AI development from the same platform.

The strategic value therefore extends beyond technical performance. By reducing architectural complexity, the platform has the potential to become increasingly embedded within enterprise data operations as organizations seek greater efficiency across expanding AI workloads.

Artificial Intelligence & Machine Learning Infrastructure

Artificial intelligence has become an increasingly important component of the Databricks platform.

The company has expanded its capabilities beyond traditional machine learning to support generative AI development, retrieval-augmented generation (RAG), model governance, vector search, AI agents, and production-scale deployment of enterprise AI applications.

Rather than competing directly with frontier model developers, Databricks positions itself as the infrastructure enabling organizations to build AI applications using their own proprietary data while maintaining governance, security, and compliance requirements.

As enterprise AI adoption accelerates, infrastructure supporting model deployment, monitoring, governance, and enterprise integration may become increasingly valuable relative to the underlying foundation models themselves.

Open Architecture & Ecosystem Integration

A defining characteristic of the Databricks platform is its commitment to open standards and interoperability.

The company continues to support widely adopted open-source technologies including Apache Spark, Delta Lake, and MLflow while integrating with leading cloud providers, enterprise software platforms, and AI model providers.

Enterprise Governance & Data Intelligence

Governance has become an increasingly important component of enterprise AI adoption.

Organizations deploying artificial intelligence at scale must manage data quality, regulatory compliance, security, permissions, auditability, and responsible AI practices across increasingly complex technology environments.

Databricks has invested significantly in governance capabilities through products such as Unity Catalog, enabling centralized management of data assets, access controls, lineage, and compliance across enterprise environments.

Strategic Platform Positioning

Databricks increasingly competes as enterprise infrastructure rather than application software.

The platform is designed to support multiple organizational functions including data engineering, analytics, machine learning, software development, governance, and AI deployment. As customers consolidate more workloads onto the platform, switching costs may increase through data integration, workflow dependency, governance configuration, and organizational standardization.

Unlike software vendors focused on individual business applications, Databricks seeks to become part of the foundational technology architecture supporting enterprise AI initiatives across the organization.

The long-term investment thesis depends less on individual product features and more on whether Databricks becomes the environment where enterprises repeatedly build, govern, and scale strategic AI workloads. Product breadth matters only if it translates into operational dependency.

Public customer case studies indicate that Databricks increasingly supports production workflows including fraud detection, demand forecasting, industrial asset monitoring, customer operations, and healthcare analytics. As these workloads become embedded within daily operations, platform replacement becomes materially more complex, increasing practical switching costs.

Product Risks & Open Questions

Several strategic questions remain central to the platform's long-term competitive position.

Key product risks include:

- Increasing feature convergence across competing cloud platforms
- Competitive expansion by hyperscalers integrating similar AI capabilities
- The pace at which enterprises adopt production AI applications
- Potential commoditization of analytics and machine learning infrastructure
- Complexity associated with large-scale enterprise implementation
- Balancing platform breadth with continued product innovation
- Maintaining interoperability while continuing to build proprietary capabilities that extend beyond its open-source foundation.
- Sustaining differentiation as enterprise AI technologies mature

The principal strategic question is whether Databricks can continue evolving faster than competing ecosystems while maintaining its position as a neutral enterprise AI platform.

Infrastructure vs. Application Platform Scorecard

Dimension	Evidence Today	Confidence
Enterprise Adoption	Strong	High
AI Platform Leadership	Strong	High
Data Infrastructure	Strong	High
Governance Capabilities	Strong	High
Ecosystem Integration	Strong	High
Switching Costs	Developing	Medium-High
Competitive Defensibility	Strong but Evolving	Medium
Infrastructure Status	Established Enterprise Platform	High

Product Assessment

Databricks has evolved beyond its origins as a data analytics platform into one of the most comprehensive enterprise AI infrastructure providers in the market. Its combination of unified data management, artificial intelligence capabilities, governance, and cloud interoperability positions the company as foundational infrastructure supporting enterprise AI deployment rather than a collection of individual software tools.

The platform's strongest competitive advantage lies in its ability to integrate multiple enterprise data and AI functions within a single operating environment. As organizations continue consolidating data infrastructure and expanding AI initiatives, this integrated approach may strengthen customer retention, increase platform dependency, and support long-term revenue expansion.

However, sustained leadership will depend on the company's ability to maintain product innovation while competing against increasingly capable cloud providers, enterprise software vendors, and AI infrastructure platforms. The next phase of diligence should therefore focus on validating customer dependency, workload expansion, switching costs, and evidence that Databricks is becoming mission-critical enterprise infrastructure rather than simply another analytics platform.

COMPETITIVE LANDSCAPE

Enterprise Data & AI Platform Competitors

Snowflake

Snowflake represents Databricks' closest strategic competitor. Both companies seek to become the primary platform for enterprise data management and artificial intelligence, although they have historically approached the market from different directions.

Microsoft Fabric

Through Azure, Microsoft Fabric, Power BI, Synapse Analytics, Microsoft 365 Copilot, and its partnership with OpenAI, Microsoft increasingly offers organizations an integrated enterprise AI ecosystem spanning productivity software, cloud infrastructure, analytics, governance, and AI development.

Amazon Web Services

Amazon Web Services continues expanding its own analytics, machine learning, and generative AI offerings through services such as SageMaker, Bedrock, Redshift, and Amazon EMR.

AWS benefits from one of the world's largest cloud infrastructure platforms and extensive enterprise adoption. Customers already operating substantial workloads within AWS may increasingly evaluate native AI and analytics services before adopting third-party platforms.

Google Cloud

Google Cloud has invested heavily in artificial intelligence, data analytics, and machine learning infrastructure through BigQuery, Vertex AI, Gemini, and its broader AI research capabilities.

Google possesses exceptional technical expertise in AI and large-scale data processing. However, enterprise adoption remains smaller than that of Microsoft and AWS in many sectors.

Traditional Enterprise Data Platforms

Oracle

Oracle remains an important competitor within enterprise databases, analytics, and cloud infrastructure.

Its extensive installed customer base, particularly among large enterprises and regulated industries, provides meaningful opportunities to extend AI capabilities across existing database environments.

IBM

While IBM's market position has evolved over recent years, its strong enterprise relationships and consulting organization remain relevant, particularly for large digital transformation initiatives.

Public evidence suggests Databricks consistently wins enterprise opportunities involving complex data engineering, machine learning, streaming, and multi-cloud environments. Conversely, competitive losses appear more common where organizations prioritize simplified analytics, native cloud integration, or lower implementation complexity.

Open-Source & AI Infrastructure Ecosystem

Beyond commercial software vendors, Databricks also competes within a broader ecosystem of open-source technologies and AI infrastructure providers.

Open-source frameworks, specialized vector databases, orchestration platforms, and emerging AI infrastructure companies continue introducing new capabilities that may reduce barriers to enterprise AI adoption.



At the same time, Databricks benefits from its own open-source heritage through Apache Spark, Delta Lake, and MLflow, helping maintain strong credibility within developer communities while encouraging ecosystem participation.

While Databricks benefits significantly from its open-source heritage, this also creates an important strategic consideration. Core technologies including Apache Spark, Delta Lake, and MLflow are widely available across competing cloud ecosystems. Databricks must therefore continuously build proprietary capabilities on top of its open-source foundation to prevent commoditization and maintain competitive differentiation through enterprise workflows, governance, AI capabilities, and operational integration rather than ownership of the underlying technologies.

Key Competitive Risks

Despite its strong market position, Databricks faces several meaningful strategic risks.

Platform Convergence

Enterprise software markets continue converging as cloud providers, database vendors, analytics companies, and AI platforms increasingly offer overlapping capabilities.

This convergence may reduce product differentiation over time while increasing pricing pressure across enterprise AI infrastructure.

Cloud Hyperscaler Competition

Microsoft, Amazon Web Services, and Google Cloud possess substantial financial resources, global infrastructure, and existing enterprise customer relationships.

Their ability to integrate AI capabilities into broader cloud offerings represents a significant long-term competitive challenge.

Pricing Pressure

As enterprise AI capabilities become more widely available, organizations may increasingly compare competing platforms based on overall ecosystem value rather than individual product functionality.

This could create pressure on pricing, contract negotiations, and long-term revenue growth.

Technological Disruption

Rapid advances in artificial intelligence, data architectures, and cloud technologies require continuous product innovation.

Companies unable to adapt quickly may experience declining competitive relevance despite current market leadership.

Competitive Assessment

Competition within enterprise AI infrastructure is likely to remain intense throughout the coming decade.

However, Databricks appears to possess several structural advantages that support its competitive position, including a unified platform architecture, deep technical expertise, strong developer adoption, extensive enterprise deployments, and broad cloud interoperability.

Unlike vendors primarily competing through proprietary cloud ecosystems, Databricks has positioned itself as a relatively neutral enterprise platform capable of operating across multiple cloud providers while integrating with leading AI technologies.

The primary competitive question is not whether Databricks has capable competitors; it clearly does. The investment question is whether customers continue consolidating increasingly mission-critical workloads onto Databricks despite bundled alternatives from hyperscalers. If customers view Databricks as a neutral, best-in-class AI data platform, the company can sustain differentiation. If customers view the platform as functionally comparable to native cloud offerings, pricing power and valuation support become more vulnerable. The commercial durability of this positioning will depend on whether enterprises continue valuing cross-cloud flexibility despite increasing incentives to consolidate workloads within native hyperscaler ecosystems.

At present, Databricks appears well positioned within one of the most competitive but strategically important segments of enterprise software. The next phase of diligence should therefore focus on customer retention, workload expansion, pricing durability, competitive win rates, and evidence that the platform continues gaining strategic importance within enterprise technology environments.

BUSINESS MODEL

Business Model Overview

Databricks operates as an enterprise software platform providing cloud-native infrastructure for data engineering, analytics, machine learning, governance, and artificial intelligence. The company primarily generates recurring revenue through enterprise software subscriptions and consumption-based pricing, enabling organizations to process, analyze, and operationalize large-scale data and AI workloads.

Unlike traditional software vendors focused on individual applications, Databricks seeks to become foundational enterprise infrastructure supporting multiple business functions. As customers consolidate data engineering, analytics, governance, and AI development within a single platform, the company has the opportunity to expand both platform usage and long-term customer value.

The business model benefits from increasing enterprise investment in artificial intelligence, cloud modernization, and digital transformation, positioning Databricks to participate across multiple areas of enterprise technology spending rather than a single software category.

Enterprise Software Characteristics

Databricks exhibits many characteristics associated with high-quality enterprise software businesses.

The platform serves large organizations operating mission-critical data infrastructure, resulting in relatively long customer relationships, significant implementation investment, and recurring revenue opportunities. Enterprise deployments frequently involve multiple business units, extensive data migration, governance implementation, and integration across existing technology environments.

These characteristics support larger contract values and create opportunities for continued expansion as customers increase platform adoption over time.

As organizations expand their use of artificial intelligence, additional workloads are increasingly deployed on existing data infrastructure, potentially strengthening customer relationships while increasing recurring revenue.

Consumption-Based Revenue Model

Databricks combines enterprise subscriptions with usage-based pricing linked to platform consumption.

Rather than relying solely on user licensing, revenue increasingly reflects the volume of computing resources, data processing, machine learning workloads, and AI applications executed through the platform. As customers increase analytical activity and deploy additional AI workloads, platform consumption has the potential to grow alongside customer value.

This model aligns commercial performance with customer adoption. Organizations deriving greater value from the platform naturally generate higher infrastructure utilization, creating opportunities for revenue expansion without requiring entirely new customer acquisitions.

However, consumption-based pricing may also introduce greater revenue variability during periods of reduced customer activity or slower enterprise technology spending.

Consumption-based pricing also increases sensitivity to customer workload optimization and broader macroeconomic conditions. Long-term revenue durability therefore depends not only on customer retention but also on continued expansion of production workloads across the installed customer base.



Consumption-based models can also amplify macroeconomic volatility. Enterprise customers can often reduce platform utilization more quickly than traditional subscription spending through workload optimization, making long-term revenue durability more dependent on continued customer expansion than contractual commitments alone.

Customer Expansion & Platform Adoption

One of the strongest aspects of Databricks' commercial model is its ability to expand within existing enterprise accounts.

Organizations often begin by deploying the platform for specific analytical or data engineering workloads before extending usage across machine learning, governance, streaming analytics, business intelligence, and enterprise AI initiatives.

As departments standardize on the platform, additional users, workloads, datasets, and business processes may be consolidated into the same operating environment.

The strongest indicator of long-term business quality is therefore not simply customer acquisition, but continued workload expansion and deeper operational integration within existing enterprise customers. Expansion driven by existing customers is generally a stronger indicator of commercial durability than growth driven primarily through new customer acquisition.

A key commercial assumption requiring further validation is whether enterprise customers consistently achieve measurable business outcomes following implementation. If implementation complexity, organizational capability, or total cost of ownership delay value realization, customer expansion and long-term platform adoption could prove weaker than expected despite strong initial customer acquisition.

Revenue Durability & Switching Costs

The long-term attractiveness of the Databricks business model depends on whether enterprise customers become operationally dependent on the platform.

Successful implementations often involve significant investment in data migration, governance frameworks, application development, workflow integration, and employee training. As organizations centralize increasing amounts of enterprise data and AI infrastructure within Databricks, switching to an alternative platform may require substantial technical, operational, and financial effort.

This creates the potential for durable customer relationships supported by expanding platform usage rather than recurring software subscriptions alone.

The key commercial question is whether customers increasingly organize critical business operations around the platform, making replacement progressively more disruptive over time.

Enterprise Sales & Commercial Strategy

Databricks primarily targets large enterprises with complex data environments and significant AI investment requirements.

The company's commercial strategy emphasizes long-term strategic relationships rather than transactional software sales. Enterprise engagements frequently involve executive sponsorship, multi-year digital transformation initiatives, technical consulting, cloud migration, and ongoing platform expansion.



Strategic partnerships with Microsoft, Amazon Web Services, Google Cloud, and global systems integrators further strengthen enterprise adoption by embedding Databricks within broader technology transformation programs.

This enterprise-focused approach supports larger contract opportunities while creating additional expansion potential across business units and geographic regions.

Infrastructure Costs & Operating Leverage

Operating a large-scale enterprise AI platform requires substantial investment in cloud infrastructure, engineering, research and development, customer support, and product innovation.

Artificial intelligence workloads, real-time data processing, model training, and inference continue increasing computational requirements across enterprise environments. AI-related compute requirements may place structural pressure on gross margins unless pricing, infrastructure efficiency, and product mix evolve alongside customer adoption.

At the same time, enterprise software businesses benefit from significant operating leverage once platforms reach scale. Continued revenue growth, increasing customer utilization, and expanding enterprise adoption may improve long-term operating margins provided infrastructure efficiency continues improving alongside platform usage.

The long-term margin profile will therefore depend on Databricks' ability to balance increasing AI infrastructure costs with productivity gains, pricing discipline, and software scalability.

Pricing Power & Commercial Risks

Several factors support Databricks' long-term pricing power.

The platform combines multiple enterprise capabilities within a unified operating environment, reducing technology complexity while supporting mission-critical business operations. Organizations that standardize on the platform may become less sensitive to incremental pricing changes than customers using isolated software tools.

Nevertheless, pricing pressure remains an important commercial consideration.

Competition from hyperscalers, enterprise software vendors, and alternative AI platforms may reduce pricing flexibility as competing technologies become increasingly comparable. Customers may also seek to optimize infrastructure spending as AI deployments mature and enterprise budgets become more disciplined.

Maintaining pricing power will depend on Databricks continuing to deliver measurable business value through platform innovation, operational efficiency, and expanding enterprise capabilities.

Business Model Assessment

Databricks' business model is attractive because it combines recurring enterprise relationships with consumption-based expansion. The strongest commercial feature is the potential for existing customers to increase usage as additional data, analytics, governance, and AI workloads are consolidated onto the platform.

Publicly available evidence indicating Net Revenue Retention above 140% supports the view that expansion within existing accounts is a major growth driver. However, Gross Revenue Retention, customer churn by segment, cohort-level expansion, workload profitability, and AI-related infrastructure economics remain undisclosed and should be central priorities for Commercial Due Diligence.



The key question is whether customer expansion reflects durable production workloads and measurable commercial value rather than temporary AI experimentation, short-term consumption spikes, or unusually elevated AI spending during the current adoption cycle.



KEY STRENGTHS, RISKS & DILIGENCE PRIORITIES

Key Strengths

Strong Position Within Enterprise AI Infrastructure

Databricks has established itself as one of the leading enterprise AI and data infrastructure platforms globally. Rather than competing within a single software category, the company participates across data engineering, analytics, governance, machine learning, and generative AI, allowing it to capture a broader share of enterprise technology spending.

Integrated Platform Architecture

One of Databricks' strongest competitive advantages is the breadth of its unified platform.

Organizations increasingly seek to consolidate fragmented data environments into a single operating platform capable of supporting analytics, governance, machine learning, and AI development. Databricks addresses this requirement through its Lakehouse architecture and Data Intelligence Platform, reducing operational complexity while enabling multiple business functions to operate from a common data foundation.

Enterprise Customer Quality

Databricks has successfully established relationships with many of the world's largest enterprises across financial services, healthcare, manufacturing, retail, telecommunications, and the public sector.

Large enterprise customers typically involve longer contract durations, higher implementation costs, and greater expansion opportunities than smaller organizations.

Strong Technology Ecosystem

The company benefits from a broad ecosystem of cloud providers, technology partners, systems integrators, and open-source communities.

Strategic relationships with Microsoft Azure, Amazon Web Services, Google Cloud, NVIDIA, Anthropic, and OpenAI strengthen Databricks' ability to integrate emerging AI technologies while maintaining broad enterprise compatibility.

Attractive Expansion Dynamics

The commercial model provides meaningful opportunities for expansion after initial deployment.

Customers frequently begin with specific analytical workloads before extending platform usage across data engineering, governance, machine learning, streaming analytics, and enterprise AI initiatives.

Exceptional Expansion Economics

Sustained Net Revenue Retention above 140% indicates that existing enterprise customers continue increasing platform adoption beyond initial deployment, reinforcing the commercial durability of the business model.

Key Risks

While several commercial and competitive risks exist, the most significant risks to the investment thesis are hyperscaler competition, long-term infrastructure economics as AI workloads expand, and the sustainability of customer expansion.

Intensifying Competitive Pressure

While Databricks currently demonstrates strong enterprise adoption, increasing competition from Microsoft Fabric, Snowflake, and native hyperscaler platforms may reduce pricing power, compress sales efficiency, and limit future workload consolidation.

Platform Convergence

Enterprise software categories are increasingly converging.

Cloud providers now offer integrated analytics, machine learning, governance, vector databases, AI agents, and foundation model services within their own ecosystems. As these capabilities become more comprehensive, customers may question whether independent platforms provide sufficient incremental value.

Valuation Expectations

Databricks has achieved one of the highest private valuations within enterprise software.

Such valuation levels imply continued rapid revenue growth, successful execution, expanding operating leverage, and sustained market leadership.

AI Infrastructure Economics

Artificial intelligence workloads require substantial computing resources, particularly as organizations increase model training, inference, vector search, and agentic AI deployment.

Although increasing platform utilization supports revenue growth, infrastructure costs may also rise as AI adoption expands.

Enterprise Sales Complexity

Large enterprise deployments typically involve lengthy procurement processes, governance reviews, cloud migration planning, implementation services, and organizational change management.

These factors may extend sales cycles, increase implementation costs, delay customer value realization, and reduce expansion opportunities within less technically mature organizations.

Diligence Priorities**Customer Expansion & Revenue Quality**

- Validate Gross Revenue Retention by customer cohort.
- Assess Net Revenue Retention sustainability across customer segments.
- Understand expansion dynamics between enterprise, mid-market, and international customers.
- Evaluate customer concentration and renewal quality.
- Determine whether AI adoption is increasing long-term platform consumption or primarily driving short-term experimentation.

Customer Retention & Expansion

Further diligence should evaluate:

- validate Net Revenue Retention drivers by separating expansion from gross churn.
- measure expansion rates across enterprise cohorts, products, and use cases.
- assess product attach rates and cross-sell adoption across AI, SQL, governance, and data engineering.
- validate whether workload expansion consistently follows initial deployment.
- determine the balance between revenue growth from existing customers versus new logo acquisition.
- validate whether customer expansion reflects sustainable production workloads rather than temporary AI experimentation, unusually elevated AI spending, or short-term consumption spikes.

Customer Adoption

Further diligence should evaluate:

- measure implementation time-to-value across different customer profiles.
- assess implementation complexity as a potential barrier to broader enterprise adoption.

- evaluate required customer capabilities, including engineering talent, AI expertise, and operational resources.
- determine whether implementation complexity affects expansion rates or customer satisfaction.
- determine what proportion of AI workloads have progressed from pilot projects into production deployments.

Platform Dependency

Further diligence should evaluate:

- determine operational dependency across customer segments.
- assess switching costs, migration complexity, and replacement timelines.
- validate whether Databricks has become embedded within mission-critical workflows.
- measure customer willingness to migrate under pricing or competitive pressure.

Competitive Position

Further diligence should evaluate:

- validate competitive win rates against Snowflake, Microsoft Fabric, BigQuery, Redshift, and native hyperscaler platforms.
- assess reasons for competitive losses across enterprise segments.
- evaluate pricing pressure resulting from bundled cloud-native offerings.
- determine whether Databricks maintains differentiation beyond technology through operational dependency and switching costs.
- determine the primary commercial reasons for competitive wins and losses, including pricing, implementation complexity, bundled cloud offerings, and product capabilities.

Financial Quality

Further diligence should evaluate:

- validate revenue quality, customer concentration, and cohort-level retention.
- assess pricing power across customer segments and enterprise renewals.
- evaluate customer willingness to absorb pricing increases relative to competitive alternatives.
- analyze infrastructure economics, including gross margins by product, AI workload profitability, cloud infrastructure costs, and operating leverage.
- assess operating leverage as AI-related compute requirements increase.
- determine the balance between expansion ARR and new customer ARR.

IPO Readiness

As Databricks moves closer to becoming a public company, diligence should focus on:

- validate sustainability of hypergrowth relative to public market expectations.
- assess readiness for public company reporting, governance, and operational discipline.
- evaluate durability of revenue quality, free cash flow generation, and operating leverage.
- determine whether AI strategy, competitive positioning, and financial profile support long-term public market valuation.
- assess whether customers view Databricks' cloud-neutral positioning as a strategic advantage or whether native hyperscaler offerings increasingly reduce the value of cross-cloud interoperability.
- assess whether hyperscaler partnerships continue supporting commercial growth or increasingly compete with Databricks through native offerings.
- determine what proportion of AI workloads have progressed from pilot projects into production deployments.

Successful execution across these areas will influence both IPO timing and long-term public market performance.

High-Conviction Milestones to Track

- Annual revenue continuing to grow above enterprise software peer averages
- Net Revenue Retention remaining above 120%
- Continued expansion within Fortune 500 customers
- Increasing adoption of AI-specific platform capabilities
- Sustained gross margin improvement despite higher AI infrastructure utilization
- Successful international expansion across major enterprise markets
- Continued growth in strategic partnerships and developer ecosystem adoption
- Clear progress toward public market readiness

Signals That Would Weaken the Investment Thesis

- Material slowdown in enterprise customer expansion
- Declining Net Revenue Retention (NRR)
- Increasing pricing pressure from hyperscalers
- Loss of technological differentiation relative to competing platforms
- Significant deterioration in infrastructure economics
- Slowing enterprise AI investment
- Customer migration toward native cloud services
- Delays to expected IPO driven by weaker operating performance

Area	High-Conviction Signal	Weakening Signal
Customer Growth	Strong expansion within existing customers	Slowing workload growth
Revenue Quality	NRR above 120%	Declining retention
Platform Adoption	Multi-product enterprise deployments	Narrow workload usage
Competitive Position	Consistent enterprise wins	Increasing losses to hyperscalers
Infrastructure	Expanding AI workload adoption	Margin compression from compute costs
Financial Performance	Strong operating leverage	Rising customer acquisition costs

Diligence Assessment

Databricks exhibits many of the characteristics associated with category-leading enterprise software businesses, including a high-quality customer base, expanding platform adoption, recurring revenue, and strong exposure to long-term AI infrastructure spending.

The investment thesis appears increasingly dependent on the company's ability to deepen enterprise platform adoption rather than simply maintain high revenue growth. The strongest evidence supporting long-term value creation would be continued expansion of customer workloads, durable switching costs, and increasing operational dependence on the platform across large enterprises.



The next phase of diligence should therefore prioritize validating customer retention, platform centrality, infrastructure economics, competitive durability, and financial quality. These factors will ultimately determine whether Databricks can sustain its leadership position as enterprise AI infrastructure evolves and competition across the market continues to intensify.

TRANSACTION LANDSCAPE & IPO CONSIDERATIONS

Enterprise AI Infrastructure Transaction Environment

Accelerating Investment Across Enterprise AI

Artificial intelligence has become one of the most active areas of investment across private capital and enterprise software. Venture capital firms, growth equity investors, sovereign wealth funds, strategic corporate investors, and public market participants continue allocating significant capital toward companies providing the infrastructure required to develop, deploy, and manage enterprise AI applications.

Unlike earlier software cycles that focused primarily on cloud migration or business process automation, the current investment environment increasingly centers on platforms capable of integrating enterprise data, governance, analytics, and artificial intelligence into unified operating environments.

This shift has supported significant valuation expansion across leading AI infrastructure companies, particularly those demonstrating strong enterprise adoption, recurring revenue growth, and strategic importance within customer technology architectures.

Recent Enterprise Software Transactions

Recent investment activity across enterprise AI infrastructure continues supporting premium valuations for category leaders.

Selected examples include:

- Databricks' Series J financing at a \$62B valuation, led by Thrive Capital and co-led by Andreessen Horowitz, DST Global, GIC, Insight Partners, and WCM Investment Management.
- Databricks' announced the acquisition of Neon, a serverless Postgres company, to strengthen its position in AI-native application and agent infrastructure.
- Google's acquisition of Wiz, reflecting strategic demand for cloud and AI security infrastructure.
- Salesforce's acquisition of Informatica, reinforcing the strategic value of enterprise data management, governance, and AI-ready data foundations.

These transactions support the view that enterprise AI infrastructure value is increasingly accruing to platforms that control data, governance, security, and operational AI deployment rather than narrow application-layer tools.

Strategic Position Within the Enterprise Software Ecosystem

Databricks occupies a strategically important position within the broader enterprise software ecosystem.

The platform sits at the intersection of cloud infrastructure, enterprise data management, analytics, machine learning, and artificial intelligence. As organizations continue consolidating these functions into integrated technology environments, infrastructure providers capable of supporting multiple enterprise workloads may become increasingly valuable.

Rather than competing solely within data analytics, Databricks increasingly participates across several of the largest enterprise software categories simultaneously.

This broad positioning increases the company's strategic relevance while reducing dependence on any single product segment.

Potential Strategic Acquirers

Although Databricks is widely expected to pursue an eventual public listing, its strategic importance makes it relevant to consider potential acquisition interest.

While a strategic acquisition cannot be ruled out, the company's scale, valuation, and institutional shareholder base make an IPO the substantially more probable exit pathway.

Hypothetical strategic acquirers could include major cloud providers, enterprise software companies, or global technology firms seeking to strengthen their positions in enterprise AI infrastructure.

IPO Considerations

Characteristics Supporting a Public Listing

Databricks demonstrates many characteristics commonly associated with successful public enterprise software companies.

These include:

- large and diversified enterprise customer base
- strong recurring revenue profile
- exposure to long-term structural growth markets
- global operations
- significant institutional investor backing
- experienced executive leadership
- expanding product portfolio
- strong ecosystem partnerships

In addition, enterprise AI remains one of the highest-priority investment themes across global capital markets, potentially providing favorable conditions for a future public listing.

Public Market Positioning

Following an IPO, Databricks would likely be evaluated alongside leading enterprise software and AI infrastructure companies rather than traditional database vendors alone.

Databricks' reported revenue run-rate above \$5.4B, growth above 65% year-over-year, positive free cash flow, and Net Revenue Retention above 140% provide meaningful IPO-readiness indicators. However, public-market investors are likely to scrutinize whether these metrics remain durable as growth scales, AI workloads increase infrastructure intensity, and competitive pressure from hyperscalers continues.

Potential public market comparisons may include:

- Snowflake
- Palantir
- MongoDB
- Confluent
- Elastic
- ServiceNow
- Microsoft (AI infrastructure exposure)

Investors are likely to focus on:

- Revenue growth
- Net Revenue Retention
- Gross margin
- Free cash flow generation
- Operating leverage
- AI platform adoption



- Competitive positioning
- Long-term market expansion opportunities

At the current private valuation, the investment case increasingly depends on Databricks demonstrating that enterprise AI adoption translates into durable revenue expansion, resilient gross margins, and operating leverage rather than temporary demand driven by the current AI investment cycle. Public market investors are likely to place greater emphasis on profitability, revenue quality, and valuation discipline than late-stage private investors.

Conditions Required for Long-Term Public Market Success

Long-term public market performance will depend on Databricks demonstrating continued leadership across enterprise AI infrastructure while successfully transitioning from a high-growth private company to a mature public software platform.

Investors will likely expect continued evidence of:

- sustained enterprise customer expansion
- strong revenue retention
- increasing operating leverage
- improving profitability
- durable competitive differentiation
- continued product innovation
- successful execution of enterprise AI strategy
- durable revenue quality supported by Gross Revenue Retention, Net Revenue Retention, and diversified enterprise customer expansion.

The ability to balance rapid innovation with operational discipline will become increasingly important following a public listing.

Bull Case

Databricks establishes itself as the dominant enterprise AI infrastructure platform underpinning data management, analytics, governance, and artificial intelligence across large organizations worldwide. Enterprise AI adoption accelerates over the coming decade, driving sustained revenue growth, expanding operating margins, and increasing customer dependency. Continued platform innovation and ecosystem expansion strengthen competitive positioning, allowing the company to emerge as one of the defining enterprise software businesses of the AI era.

Bear Case

Enterprise AI adoption continues growing, but implementation complexity, pricing pressure, and expanding native cloud capabilities reduce customer consolidation around independent platforms. Revenue growth moderates as organizations increasingly adopt bundled cloud-native alternatives rather than standalone enterprise AI infrastructure.

Exit Perspective

Databricks is well positioned to pursue a successful public market listing supported by its scale, commercial momentum, strategic positioning, and exposure to one of the fastest-growing areas of enterprise software.

Unlike many late-stage private companies, the principal investment question is unlikely to be whether an exit opportunity exists. Rather, the more important consideration is the quality and durability of long-term value creation following a public listing.



The ultimate investment outcome will depend on whether Databricks can continue expanding from a leading data platform into indispensable enterprise AI infrastructure with durable customer relationships, strong financial performance, and sustained competitive differentiation.

If successful, the company has the potential to become one of the defining public enterprise software companies of the next decade. Conversely, if competitive convergence reduces differentiation or enterprise AI adoption develops more slowly than expected, valuation expansion may become increasingly constrained despite continued business growth.

CONVICTION ASSESSMENT

Databricks has established itself as one of the leading enterprise AI infrastructure companies, combining data engineering, analytics, governance, machine learning, and artificial intelligence within a unified cloud-native platform. The company's continued product expansion, broad enterprise adoption, and strategic partnerships position it at the center of one of the most significant long-term technology transitions since the emergence of cloud computing.

The strongest aspect of the investment thesis is Databricks' positioning as foundational enterprise infrastructure rather than a standalone analytics platform. As organizations increasingly deploy artificial intelligence across mission-critical business functions, demand for platforms capable of managing proprietary data, governance, and AI development is likely to continue expanding. Databricks is well positioned to benefit from this structural shift through its integrated architecture, open ecosystem, and growing enterprise footprint.

Commercial momentum also supports the investment case. The company has attracted a high-quality enterprise customer base across multiple industries while maintaining strong institutional backing and significant financial resources. Public information suggests continued expansion in platform capabilities, customer adoption, and strategic partnerships, reinforcing its position as a leading participant within the enterprise AI ecosystem.

However, the long-term investment outcome remains dependent on continued execution. Competition across enterprise AI infrastructure is intensifying as hyperscalers, cloud providers, and established enterprise software vendors expand their own AI capabilities. Sustaining product leadership while maintaining pricing power and attractive unit economics will become increasingly important as the market matures.

The company must also demonstrate that enterprise customers continue expanding workloads and embedding critical business processes within the platform. Long-term value creation is likely to depend less on initial adoption and more on the durability of customer relationships, workload expansion, and operational dependence across increasingly complex enterprise environments.

Bull Case

Databricks becomes the dominant enterprise AI infrastructure platform, serving as the foundational operating environment for data management, analytics, governance, and artificial intelligence across global enterprises. Continued AI adoption accelerates customer expansion, strengthens switching costs, and supports sustained revenue growth alongside improving operating leverage. The company successfully transitions into a leading public enterprise software business with durable competitive advantages and long-term platform leadership.

Base Case

Databricks maintains its position as one of the leading enterprise AI infrastructure providers while operating within an increasingly competitive market. Enterprise AI adoption continues supporting healthy revenue growth, although competition from hyperscalers and cloud-native platforms moderates market share gains and pricing power. The company remains a high-quality enterprise software business capable of generating attractive long-term returns, supported by continued customer expansion and platform innovation.

Bear Case

Enterprise AI infrastructure becomes increasingly commoditized as cloud providers integrate comparable capabilities directly into broader cloud ecosystems. Customers distribute workloads across multiple competing platforms rather than consolidating around Databricks, limiting switching costs and reducing pricing power. Revenue growth slows, infrastructure costs remain elevated, and long-term valuation expectations prove difficult to sustain despite continued enterprise adoption.

Conviction Level**Medium-High Conviction**

Based on publicly available information, Databricks demonstrates many of the characteristics associated with category-leading enterprise software businesses, including a large addressable market, strong enterprise adoption, differentiated technology, recurring revenue characteristics, and exposure to long-term structural growth in artificial intelligence.

The primary factors supporting conviction include:

- leadership within enterprise AI infrastructure
- integrated platform spanning data, analytics, governance, and AI
- high-quality enterprise customer base
- strong ecosystem partnerships
- attractive long-term market dynamics
- proven product innovation and commercial execution

Areas requiring further validation include:

- Gross Revenue Retention and customer churn quality.
- sustainability of Net Revenue Retention above 140%.
- enterprise workload expansion across customer cohorts.
- implementation time-to-value across customer profiles.
- pricing durability against bundled cloud-native competitors.
- long-term switching costs beyond flagship enterprise customers.
- AI monetization beyond traditional analytics workloads.
- customer concentration and revenue diversification.
- infrastructure economics as AI workloads increase.

Investment Kill Conditions

During Commercial Due Diligence, particular attention should be given to developments that could materially weaken the investment thesis, including:

- sustained Net Revenue Retention deterioration indicating weaker long-term expansion dynamics.
- accelerating competitive displacement by hyperscaler-native platforms that materially reduce Databricks' pricing power or enterprise win rates.
- material deterioration in AI workload economics resulting in sustained gross margin compression.

Recommendation

Based on the current commercial assessment, Databricks warrants progression to full Commercial Due Diligence.

The next phase of diligence should focus on validating customer retention, workload expansion, implementation complexity and time-to-value, competitive win rates, pricing durability, infrastructure economics, and evidence that enterprise customers increasingly view Databricks as operational infrastructure rather than another analytics platform.

The key underwriting question is not whether enterprise AI spending will grow, but whether spending consolidates around a small number of strategic platforms or remains fragmented across competing ecosystems.

The principal commercial risk is not whether enterprises invest in AI infrastructure, but whether customers consistently realize sufficient commercial value to justify expanding platform adoption over time.



Commercial Due Diligence should focus on validating revenue quality, Gross Revenue Retention, customer concentration, workload expansion, implementation time-to-value, pricing durability, competitive win rates, infrastructure economics, and long-term switching costs. Particular attention should be given to whether customer adoption is translating into durable platform dependency and measurable commercial value across customer segments.

Source: Company website, publicly available information, industry reports, market research, and Cohres analysis.

Disclaimer: This assessment has been prepared using publicly available information and Cohres analysis. Conclusions represent commercial evaluation and professional judgment based on available information at the time of writing. This document is provided for informational purposes only and does not constitute investment, legal, accounting, or financial advice.

Cohres supports venture capital funds, private equity firms, family offices, hedge funds, asset managers, investment banks and growth-stage companies through commercial evaluation, commercial due diligence, investment research, execution risk assessment, conviction allocation, and AI Concierge Systems.

Learn more:

Cohres

<https://cohres.com>

Services

<https://cohres.com/services/>

AI Concierge Systems

<https://cohres.com/ai-concierge>