ISG Provider Lens®

Databricks Ecosystem Partners

Modernization and AI/ML Enablement Services

Managed Data and Optimization Services

Assessing service providers enabling enterprise Databricks transformations

QUADRANT REPORT | APRIL 2026 | GLOBAL

Customized report courtesy of:
Tiger Analytics

Report Author: Gowtham Sampath and Hemangi Patel

Databricks unifies data modernization, governance and scalable AI architecture for enterprises

Databricks has rapidly evolved into a central pillar of enterprise data and AI transformation strategies. Organizations are reevaluating the architectural core that supports data engineering, analytics, ML and GenAI workloads. The Lakehouse architecture, which unifies the scalability of data lakes with the reliability and performance of data warehouses, has emerged as a strategic response to the growing complexity of enterprise data estates. Capabilities including Delta Lake, Unity Catalog, MLflow, Mosaic AI and serverless compute position Databricks not merely as an analytics platform but as a comprehensive data-to-AI control plane. This platform evolution is unfolding against a broader backdrop of economic, technological and regulatory shifts that are reshaping enterprise data strategies globally.

Market context: data modernization in an AI-driven economy

The global market for data modernization is being shaped by macroeconomic pressure, accelerating AI adoption and structural shifts in enterprise operating models. Insights from ISG's **2025 Data and AI Programs** study indicate that despite heightened cost scrutiny and economic uncertainty, enterprises continue to expand investment in data and AI initiatives. Market sentiment toward data remains strong, with nearly three-quarters of organizations viewing data as a valuable asset worthy of sustained investment. However, this optimism is tempered by concerns around cost and security. Nearly 40 percent indicate that security risks reduce their willingness to collect and retain more data, and a similar proportion expresses concern that the cost of harmonizing and managing data across the enterprise may outweigh the benefits. This duality captures the tension at the core of modern data strategies. Enterprises recognize the immense potential of AI-driven insights, yet they remain cautious about governance exposure, compliance liabilities and operational overhead.

Databricks enables AI-first modernization with governed, scalable data foundations

Executive Summary

Data usability for AI has emerged as the leading anticipated challenge for 2025 and 2026. Fragmented silos, inconsistent standards and limited interoperability between legacy and cloud platforms continue to constrain scalable AI deployment. Organizations struggle to centralize, standardize and prepare datasets for LLM training, embedding and vectorization. Modernization is, therefore, increasingly AI-led rather than reporting-led. Enterprises are modernizing to make data AI-ready, with unified ingestion, transformation, governance and lifecycle management becoming foundational requirements.

Demonstrating measurable value further intensifies market pressure. Measuring ROI on data and AI investments remains a top concern, particularly as boards demand clearer justification for AI-related capital allocation. Benefits such as improved decision speed or risk reduction are difficult to quantify, reinforcing the need for integrated observability, cost transparency and performance tracking within modernization programs. Additionally, GenAI adoption continues to advance, though production maturity remains uneven.

Enterprises are expanding use cases cautiously while navigating integration gaps, governance complexity and cost variability. This measured progression strengthens demand for platforms capable of supporting experimentation and production within a consistent governance and financial framework.

Within this context, the Databricks ecosystem aligns closely with enterprise priorities. The Lakehouse architecture addresses fragmentation by unifying structured and unstructured data under a single control plane. Integrated governance capabilities respond to regulatory and security expectations, while native AI functionality supports AI-readiness objectives. Real-time observability and FinOps integration address growing demands for cost accountability and measurable outcomes. As enterprises seek to simplify complexity and scale AI responsibly, unified data platforms are becoming strategic enablers of sustainable transformation.

Strategic imperatives shaping the Databricks ecosystem

In light of current market dynamics, the evolution of the Databricks ecosystem aligns with three defining enterprise imperatives. First, enterprises are consolidating fragmented data estates spanning legacy warehouses, on-premises systems and multicloud environments. Second, enterprises are accelerating AI and GenAI adoption, moving from experimentation toward enterprise-wide deployment. Third, regulatory scrutiny and cost pressures are compelling enterprises to institutionalize governance, transparency and FinOps discipline across the data stack. Databricks provides an integrated foundation that addresses these shifts holistically, while service providers are shaping modernization and managed optimization strategies around these themes.

Strategic imperative 1: consolidation of fragmented data estates

- **Enterprise perspective: structural modernization imperatives**
  Enterprises continue to operate complex data landscapes built over years of incremental investments. Legacy warehouses such as Teradata, Oracle and SQL Server co-exist with siloed data marts, ETL pipelines and cloud-native analytics tools. These environments were originally optimized for structured, batch-oriented reporting rather than real-time analytics or AI-driven workloads.

  The fragmentation creates duplicated datasets, inconsistent definitions of business metrics, prolonged data preparation cycles and escalating licensing costs. Operational inefficiencies intensify when structured and unstructured data must be integrated for advanced analytics or AI use cases. Logs, streaming data, transactional systems and document repositories often remain disconnected from centralized analytics frameworks.

  Enterprises increasingly recognize that modernization must go beyond infrastructure migration. Consolidation requires establishing a single governed data backbone capable of supporting ingestion, transformation, analytics and AI from a unified architectural layer. The objective is to eliminate silos, reduce operational redundancy and enable scalable innovation on a consistent data foundation.

The convergence of transactional and analytical workloads further influences consolidation strategies. Enterprises seek low-latency data-serving capabilities that power AI-driven applications while maintaining analytical depth. The demand for near real-time intelligence necessitates architectures that bridge operational and analytical systems.

- **Service provider perspective: AI-driven lakehouse modernization**
  Service providers are responding with structured modernization frameworks centered on Databricks Lakehouse. Automation and AI-driven migration toolkits are accelerating transitions from legacy warehouses and ETL platforms while minimizing manual refactoring. Migration methodologies increasingly incorporate prebuilt connectors, reference architectures and reusable intellectual property to reduce risk and shorten time-to-value.

  Industry-specific accelerators are embedded into consolidation strategies to align transformation efforts with industry requirements. Financial services implementations emphasize regulatory reporting and data lineage, while healthcare deployments prioritize data privacy and interoperability. Manufacturing solutions integrate streaming data and predictive maintenance use cases. Verticalized templates enhance confidence and reduce implementation friction.

  Integration plays a pivotal role in modernization narratives. Strong demand exists for integrating enterprise systems such as SAP Business Data Cloud with Databricks to unify transactional and analytical insights. These integration frameworks allow enterprises to extract greater value from trusted operational systems while extending analytics and AI capabilities. Service providers increasingly co-develop and co-sell solutions with Databricks to ensure architectural alignment. Joint engineering initiatives, shared product roadmaps and coordinated GTM motions strengthen modernization outcomes. Close collaboration enables providers to incorporate emerging platform capabilities early in client engagements, reinforcing long-term ecosystem alignment.

Strategic imperative 2: acceleration of AI and GenAI at enterprise scale

- **Enterprise perspective: scaling AI beyond pilots**
  Enterprises have invested significantly in AI experimentation, yet scaling AI across business units remains a persistent challenge. Disconnected data pipelines, inconsistent MLOps practices, limited model lineage visibility and fragmented tooling slow the transition from pilot to production. AI initiatives often struggle to demonstrate repeatability and governance consistency.

  GenAI and agentic AI introduce additional complexity. Enterprises must manage vector search infrastructure, embedding pipelines, prompt governance, model evaluation and latency optimization. Human-in-the-loop validation processes are essential to ensure output quality and accountability. The transition from isolated prompt-based experimentation to domain-specific LLM deployments embedded within business workflows requires architectural cohesion and operational rigor. Organizations also face pressure to integrate AI across enterprise systems rather than confining use cases to innovation labs. Decision intelligence platforms must operate within secure, auditable frameworks while maintaining cost efficiency and performance reliability.
- **Service provider perspective: AI-native architectures and agentic enablement**
  Service providers are aligning AI enablement strategies with Databricks' AI-native capabilities. Databricks Lakehouse architectures converge data engineering, analytics, ML and GenAI into a single governed environment. Mosaic AI and MLflow are leveraged to streamline model development, lifecycle management and domain-specific LLM deployment.

  Providers are building reusable AI accelerators that incorporate vector search, prompt templates and evaluation mechanisms. These accelerators reduce experimentation cycles and accelerate production readiness. Agentic AI frameworks

are emerging within service portfolios, enabling orchestrated workflows with policy-guarded controls and secure state management.

The convergence of MLOps and LLMOps into unified lifecycle frameworks reflects a broader shift toward XOps by design. Enterprises increasingly expect FinOps, MLOps, AgentOps and observability to function cohesively rather than independently. Service providers are institutionalizing these capabilities within AI enablement programs to reduce operational overhead and improve reliability.

Additionally, co-innovation with Databricks further enhances AI acceleration efforts. Participation in early access programs and innovation labs allows providers to embed emerging AI capabilities directly into client engagements. Joint solution engineering strengthens alignment between service frameworks and platform evolution. Coordinated co-sell initiatives position AI modernization as a shared strategic narrative, increasing enterprise confidence in scalability and long-term viability.

Strategic imperative 3: governance, transparency and FinOps discipline

- **Enterprise perspective: trust, accountability and cost predictability**
  Regulatory scrutiny surrounding data privacy, fairness, explainability and compliance continues to intensify across industries. Enterprises must enforce fine-grained access control, implement PII masking policies, maintain automated lineage tracking and monitor bias across AI models. Governance has shifted from reactive compliance to proactive trust-building across the data-to-AI lifecycle.

  Cost management presents another critical dimension for enterprises adopting Databricks. AI workloads increase cloud consumption, creating variability in spending patterns. Finance and technology stakeholders require predictive scaling, Databricks Unit (DBU) tracking, workload right-sizing and measurable return on modernization investments. Migration initiatives must demonstrate improved TOC while preserving performance and scalability.

  Operational resilience is equally important. Real-time monitoring, SLA-based management and anomaly detection are essential to maintain reliability across data pipelines and AI models. Enterprises increasingly seek unified observability frameworks capable of providing end-to-end visibility into performance, cost and governance metrics. Secure collaboration further expands governance requirements. Cross-enterprise analytics initiatives demand trusted data sharing mechanisms that preserve compliance while enabling innovation.

- **Service provider perspective: managed optimization and continuous governance**
  Service providers are increasingly repositioning managed services from reactive support models to proactive optimization partnerships. Governance and cost management are no longer treated as downstream compliance activities; instead, they are embedded into the core architecture and operating model from the outset. Providers are standardizing on Unity Catalog as the central governance layer across client environments, enabling consistent policy management, lineage transparency, and access control across data and AI assets.

  FinOps maturity is becoming a visible differentiator across managed Databricks engagements. Consumption intelligence, workload right-sizing strategies and cluster performance optimization are being embedded directly into operational playbooks. Engagement models increasingly include structured cost forecasting, budget guardrails and usage accountability frameworks that align engineering decisions with financial objectives.

  Real-time observability is emerging as a defining feature of advanced managed services offerings within the ecosystem. Service providers are developing unified monitoring dashboards that combine cost consumption trends, SLA adherence metrics and data quality signals into integrated control views. These dashboards move beyond passive visualization and incorporate intelligent alerting, anomaly detection and predictive optimization recommendations.

Observability is increasingly positioned as an operational command center that enables early issue detection, faster remediation and continuous performance tuning.

Secure collaboration is also becoming a structured component of managed optimization strategies. Providers are incorporating Delta Sharing and clean-room architectures into enterprise data strategies to enable governed, multi-party analytics. This capability supports ecosystem-based innovation while maintaining strict compliance controls and auditability standards.

Overall, managed optimization within the Databricks ecosystem is evolving, and service providers are aligning governance, FinOps and observability into a unified operating model that supports sustained AI scale, cost discipline and operational resilience rather than periodic post-implementation adjustments.

Conclusion and market outlook

The Databricks ecosystem is evolving from a modernization platform into a strategic foundation for AI-driven enterprise transformation. As organizations confront fragmented estates, expanding AI ambition, regulatory complexity and cost scrutiny, the focus is shifting toward scalable operationalization. Data consolidation will remain central for enterprises that seek to scale AI and build a governed data foundation capable of real-time decision intelligence. Enterprises that fail to unify data estates under governed architectures will struggle to deploy agentic systems and real-time decision intelligence. Strategic emphasis must move from incremental integration to enterprise-wide convergence anchored in unified platforms.

AI adoption is also entering a more disciplined stage. Production-grade, domain-aligned AI embedded within workflows will replace isolated pilots. Enterprises will demand measurable ROI, governance assurance and end-to-end transparency. Architectures integrating lifecycle management, observability and cost optimization will define AI maturity. Governance and cost discipline will shape differentiation. Regulatory mandates and cybersecurity risks elevate governance into a strategic capability, while macroeconomic pressure reinforces cost visibility and predictive scaling. Real-time observability and optimization will become standard components of managed environments.

The partner ecosystem will be critical. Providers capable of delivering unified modernization-to-optimization lifecycle models will accelerate enterprise outcomes. Over the next three to five years, competitive advantage will depend less on tool selection and more on architectural discipline and ecosystem alignment. Enterprises that treat Lakehouse as a strategic control plane for the data-to-AI value chain will scale innovation while managing risk. Those that combine modernization with continuous optimization and governed AI execution will define the next era of enterprise performance.

Databricks is becoming the strategic lakehouse foundation that enables enterprises to modernize fragmented data estates, embed governance and FinOps by design, and scale AI and agentic innovation through unified, continuously optimized data-to-AI operations.

This study evaluates providers' capabilities based on their **service portfolio and competitiveness** within the Databricks ecosystem.

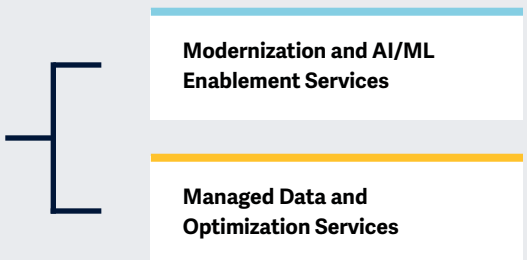

Simplified Illustration Source: ISG 2026

Definition

Databricks has become a pivotal platform driving the convergence of data engineering, analytics and AI. Its Unified Data Analytics Platform integrates data, AI and governance to accelerate time to insight and enable enterprise-scale GenAI. Built on the lakehouse architecture, Databricks unifies structured, unstructured and streaming data under a single governance framework powered by Delta Lake, Unity Catalog, MLflow and Mosaic AI. The introduction of MosaicML and Databricks AI/ agent frameworks strengthens its role as anenabler of enterprise-grade GenAI and custom LLM development, where governed data meets scalable AI capabilities.

The Databricks partner ecosystem has expanded rapidly, reflecting strong enterprise demand for modernization, governance and AI-ready platforms. Partners (providers) are critical in translating Databricks' vision into tangible business outcomes, spanning strategy, migration, operational optimization and innovation. As Databricks continues to natively integrate data intelligence and AI, it is evolving from a unified data management solution into a strategic AI backbone.

Providers' capabilities broadly fall into two quadrants:

- **Modernization and AI/ML Enablement Services**, including advisory, modernization, data engineering, data sharing and collaboration, and business application development.
- **Managed Data and Optimization Services**, including post-implementation support and cost optimization, with a focus on security, regulatory compliance and observability.

ISG's analysis evaluates service providers delivering Databricks-aligned consulting, implementation and managed services that enable scalable, secure, and cost-efficient data and AI transformation.

Scope of the Report

This ISG Provider Lens® quadrant report covers the following two quadrants for services/solutions: Modernization and AI/ML Enablement Services and Managed Data and Optimization Services.

This ISG Provider Lens® study offers IT decision-makers:

- Transparency on the strengths and weaknesses of relevant providers
- A differentiated positioning of providers by segments (quadrants) based on their competitive strengths and portfolio attractiveness
- Focus on the Global and Brazil markets

Our study serves as the basis for important decision-making by covering providers' positioning, key relationships and go-to-market considerations. ISG advisors and enterprise clients also use information from these reports to evaluate their existing vendor relationships and potential engagements.

Provider Classifications

The provider position reflects the suitability of providers for a defined market segment (quadrant). Without further additions, the position always applies to all company sizes classes and industries. In case the service requirements from enterprise customers differ and the spectrum of providers operating in the local market is sufficiently wide, a further differentiation of the providers by performance is made according to the target group for products and services. In doing so, ISG either considers the industry requirements or the number of employees, as well as the corporate structures of customers and positions providers according to their focus area. As a result, ISG differentiates them, if necessary, into two client target groups that are defined as follows:

- **Midmarket:** Companies with 100 to 4,999 employees or revenues between $20 million and $999 million with central headquarters in the respective country, usually privately owned.
- **Large Accounts:** Multinational companies with more than 5,000 employees or revenue above $1 billion, with activities worldwide and globally distributed decision-making structures.

The ISG Provider Lens® quadrants are created using an evaluation matrix containing four segments (Leader, Product & Market Challenger and Contender), and the providers are positioned accordingly. Each ISG Provider Lens® quadrant may include a service provider(s) which ISG believes has strong potential to move into the Leader quadrant. This type of provider can be classified as a Rising Star.

- **Number of providers in each quadrant:** ISG rates and positions the most relevant providers according to the scope of the report for each quadrant and limits the maximum of providers per quadrant to 25 (exceptions are possible).

Modernization and AI/ML Enablement Services

Who Should Read This Section

This report is valuable for service providers offering Databricks modernization and AI/ML enablement services globally to understand their market position and for enterprises looking to evaluate these providers. In this quadrant, ISG highlights the current market positioning of these providers based on the depth of their service offerings and market presence.

Technology professionals

should read this report to assess Databricks implementation partners' expertise in data engineering, governance, cloud architecture and multicloud deployment strategies. The report highlights providers skilled in Databricks native engineering, secure data sharing, multicloud architecture and MLOps, helping them accelerate migration, reduce delivery risk and operationalize AI at scale.

Business and strategy professionals

should read this report to track evolving Databricks-related trends, assess providers' ability to drive enterprise data maturity and cloud readiness, and identify strategic use cases. The report offers insights into advisory strength, modernization roadmaps and AI/ML enablement strategies that influence long-term business outcomes.

Procurement professionals

should read this report to evaluate providers offering strategic advisory and end-to-end delivery for Databricks implementations. The report guides the selection of partners that deliver measurable cost optimization and compliance assurance. It provides insights into vendor capabilities, contractual flexibility and risk mitigation strategies to support informed sourcing decisions aligned with enterprise transformation goals.

Research and innovation professionals

should read this report to assess providers' ability to harness AI, ML and Databricks' data capabilities for innovative, data-driven solutions. The report helps identify leaders driving next-generation technologies and advanced analytics, while highlighting partner roadmaps that enable cutting-edge solutions and future-proof enterprise strategies.

ISG Provider Lens™

Source: ISG RESEARCH

Databricks Ecosystem Partners
Modernization and AI/ML Enablement Services

Global 2026

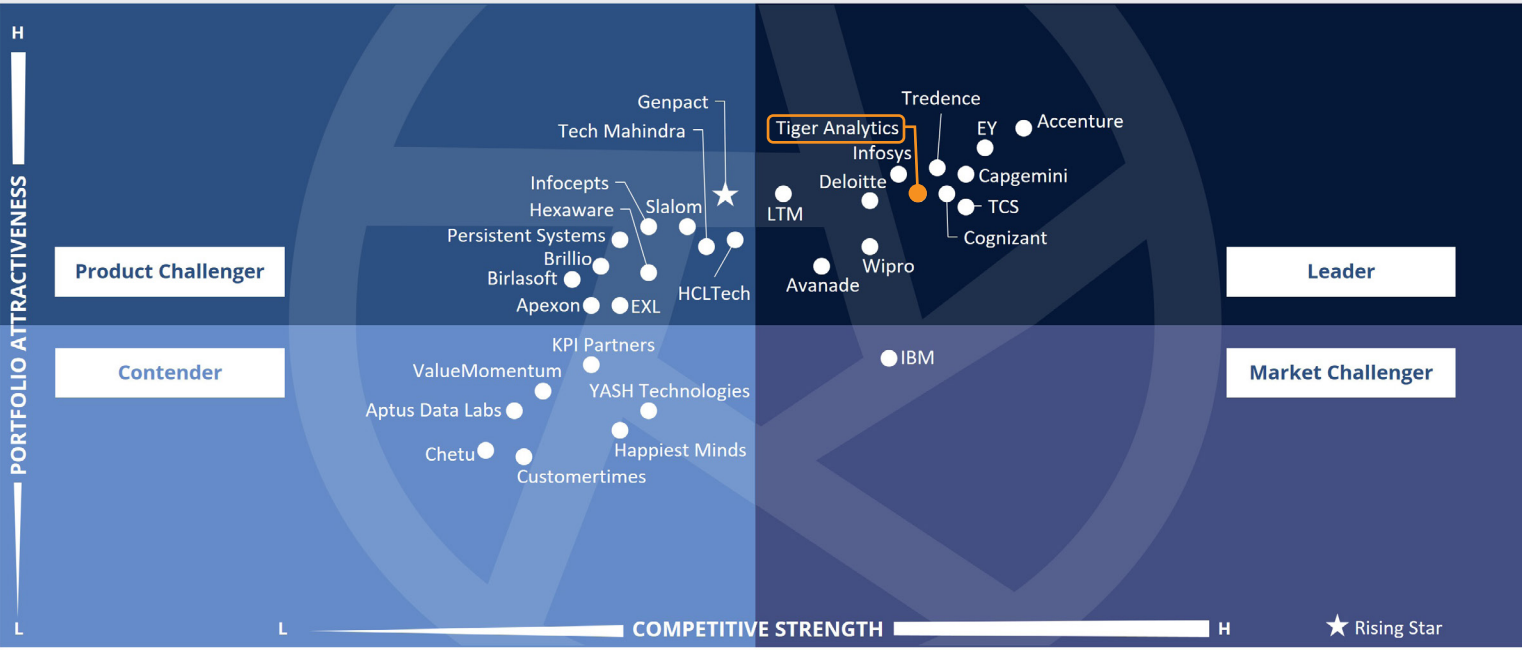

This quadrant assesses providers delivering **end-to-end Databricks modernization**, spanning **strategy, architecture, migration, and deployment** to enable scalable data and AI transformation, with strengths in **Lakehouse implementation, real-time analytics, and AI operationalization.**

Gowtham Sampath and Hemangi Patel

Definition

This quadrant evaluates providers that deliver end-to-end advisory, architecture, migration and deployment services on Databricks, helping enterprises define data and AI modernization strategies, implement Databricks Lakehouse Platform and accelerate AI adoption. Typical engagements include cloud and data modernization road maps, architecture and migration planning, data engineering, analytics enablement and governance setup through Unity Catalog.

Providers are evaluated on expertise in data architecture, Delta Lake and MLflow implementation, and multicloud or hybrid data integration, as well as their ability to operationalize streaming pipelines, enable real-time analytics and incorporate Mosaic AI into enterprise workflows. Leading providers combine strategic consulting with strong technical execution and reusable IP assets to shorten time to value.

Eligibility Criteria

1. Ability to offer **end-to-end Databricks services**, encompassing strategy, consulting, architecture design, migration planning and implementation
2. Expertise in **developing analytics frameworks, architectural blueprints, adoption road maps, AI and ML workflows**, and industry-specific data solutions and applications
3. Proven experience in **Databricks-based data modernization, migration and AI enablement** projects
4. Strong capabilities in **designing lakehouse and multicloud data architectures** using Delta Lake, Unity Catalog and MLflow
5. Referenceable **implementations involving analytics enablement, streaming data, or AI and ML solutions** using Databricks components
6. Documented **frameworks, accelerators or toolkits for ETL/ELT, governance or model lifecycle management**
7. Established **best practices for Unity Catalog setup, data security and lineage configuration**
8. **Capabilities in change management, user adoption** and enterprise enablement programs
9. Availability **of validated client case studies** or references showcasing Databricks-enabled outcomes
10. **Ecosystem alignment and strategic partnerships, including preferred partner designations or early-access participation**, that demonstrate innovation maturity

Leader

"Tiger Analytics demonstrates strong Databricks enablement, sustaining modernization, analytics consumption and AI initiatives through robust governance and cost discipline, with repeatable solution constructs aligned to Databricks."

Gowtham Sampath and Hemangi Patel

Tiger Analytics

Overview

Tiger Analytics is headquartered in California, U.S. It employs over 6,000 associates across offices in North America, Latin America, Europe and APAC. The company aligns Databricks modernization initiatives to defined business outcome maps, tying lakehouse transformation directly to enterprise value. Through industry-specific accelerators and structured co-innovation programs, it connects data foundation buildout with AI and ML enablement, reinforcing repeatable, domain-driven modernization pathways across priority sectors.

Strengths

Outcome map alignment: Tiger Analytics aligns modernization and AI use cases to Databricks outcome maps, linking data foundation buildout to defined business value themes, such as customer, workforce and supply chain priorities. This approach connects platform engineering to structured value pathways. By mapping accelerators to Databricks outcome constructs, Tiger Analytics reinforces repeatability and supports joint field alignment across modernization and AI enablement programs.

Industry-focused AI enablement: Tiger Analytics embeds AI and ML capabilities directly into modernization journeys through domain-specific accelerators developed on Databricks. Solutions such as Always On Brand Equity, Next Best Action and AI-driven Model Risk Management demonstrate the company's ability to translate platform capabilities into operational intelligence across retail, consumer goods, financial services, life sciences and manufacturing. These solutions reflect the company's maturity in linking modernization efforts to business outcomes.

Co-innovation model: Tiger Analytics participates in joint solution initiatives and selective Databricks programs, contributing to industry solution incubation and coordinated GTM execution. Its involvement in co-developed workshops and solution development reflects alignment with Databricks' evolving roadmap and supports AI-driven modernization at scale.

Caution

Tiger Analytics could further institutionalize its Databricks modernization frameworks into repeatable industry blueprints with quantified outcome benchmarks, strengthening differentiation as enterprise programs scale across regions and multicloud environments.

Managed Data and Optimization Services

Who Should Read This Section

This report is valuable for service providers offering Databricks managed data and optimization services globally to understand their market position and for enterprises looking to evaluate these providers. In this quadrant, ISG highlights the current market positioning of these providers based on the depth of their service offerings and market presence.

Technology professionals

should read this report to understand how providers deliver continuous platform optimization and AI-driven monitoring. The report equips them with strategies for secure governance and operational efficiency to sustain long-term performance in the Databricks environment. It describes approaches to enhance query performance, reduce costs and simplify workload management, while highlighting integration practices across analytics, ML and business intelligence tools.

Data management and governance professionals

should read this report to explore strategies for maintaining data integrity and compliance while optimizing Databricks performance. They will discover how managed services and AI-driven monitoring ensure secure, reliable data operations. The report explains how providers address data security and privacy needs across their Databricks managed and support services portfolio.

Procurement professionals

should read this report to compare Databricks' managed and support service providers globally. The report helps identify vendors that deliver proactive value realization and cost control and transform maintenance into a strategic function for measurable improvements. The report shows how MSPs deliver continuous optimization, AI-driven monitoring and secure governance, directly influencing vendor selection and performance management.

Field service professionals

should read this report to learn how providers ensure operational continuity and efficiency through AI-driven maintenance and release management. The report offers practical insights for adopting data-driven practices that reduce downtime, improve service delivery, optimize scheduling and enhance customer satisfaction, while positioning organizations to scale efficiently in the Databricks ecosystem.

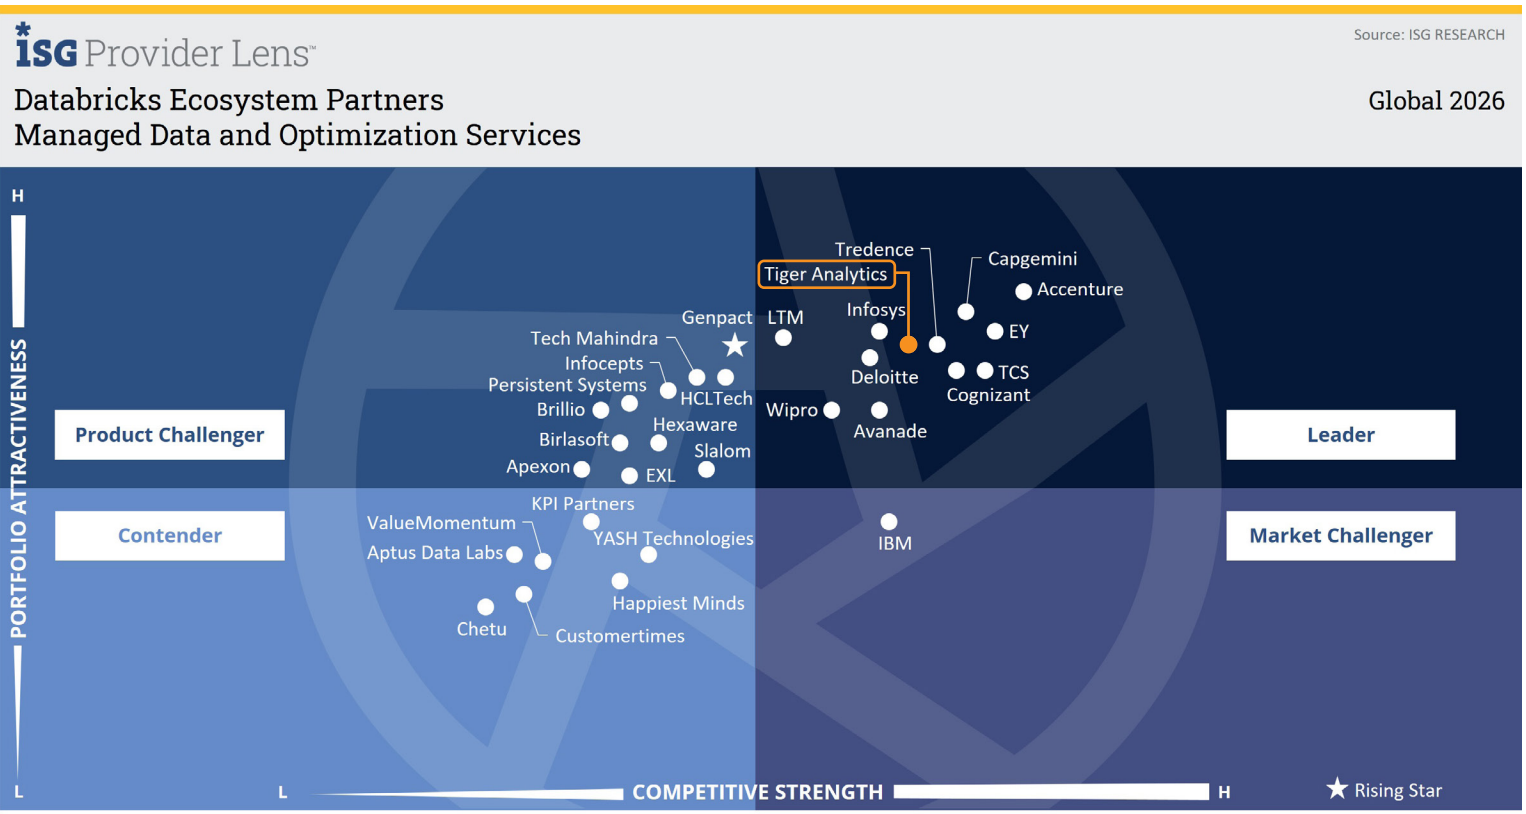

This quadrant assesses providers delivering **post-deployment Databricks operations**, focusing on **optimization, governance, and FinOps** to ensure scalable, secure environments, with strengths in **automation, observability, and cost-performance management**.

Gowtham Sampath and Hemangi Patel

Definition

This quadrant assesses providers that deliver post deployment operations, optimization and governance for Databricks environments. These managed services span monitoring, cost management, FinOps, performance tuning, compliance and MLOps automation, ensuring the reliability, scalability and regulatory adherence of Databricks workloads across hybrid cloud and multicloud infrastructures.

In this quadrant, providers are evaluated on their ability to enable intelligent operations through automation, observability and policy-driven governance. Leading providers demonstrate depth in performance optimization, cost forecasting, lineage and audit management, and secure access control. Providers are increasingly combining FinOps, MLOps and AI-driven observability into unified management frameworks that sustain operational excellence and maximize ROI.

Eligibility Criteria

1. Capabilities in delivering **Databricks-specific managed services**, including **monitoring, FinOps, optimization, SLA management, incident management** and platform operations support
2. Established **CI/CD automation for Databricks workflows**, with seamless model lifecycle governance, drift detection and rollback support
3. Expertise in **MLOps and LLMOps automation**, including model retraining, evaluation and deployment workflows
4. Expertise in **implementing enterprise-grade access controls, data protection and continuous compliance monitoring**
5. Expertise in **implementing Unity Catalog for access control, lineage tracking and regulatory compliance**
6. Ability to enable **continuous data quality monitoring** and maintain **audit-ready governance practices**
7. Ability to deliver **AI-driven observability, anomaly detection, automated remediation, proven FinOps-driven cost optimization** and centralized monitoring dashboards for platform efficiency
8. Documented success in improving **platform efficiency, reliability and cost transparency** through continuous performance tuning and optimization
9. Ability to **offer user training, certification programs and continuous environment optimization**
10. Customer success stories showcasing **measurable outcomes for managed services on Databricks**

Leader

"Tiger Analytics applies disciplined Databricks managed optimization by integrating structured cost governance, serverless adoption and operational triage into a sustained post-deployment ownership model that supports efficient platform evolution."

Gowtham Sampath and Hemangi Patel

Tiger Analytics

Overview

Tiger Analytics is headquartered in California, U.S. It employs over 6,000 associates across offices in North America, Latin America, Europe and APAC. The company embeds structured DBU governance, Unity Catalog enforcement and lifecycle controls into managed Databricks services. Through cluster optimization, AI-driven triage workflows and ML and GenAI monitoring frameworks, it supports cost transparency across data and AI environments. This approach positions managed services as a continuous optimization layer that balances consumption discipline, compliance enforcement and sustained performance across AI ecosystems.

Strengths

Structured DBU governance: Tiger Analytics embeds DBU optimization into its managed services model. Cluster sizing, configuration alignment and storage usage are reviewed through structured assessment cycles. Its remediation actions reduce unnecessary consumption while sustaining stability. The company's serverless guidance supports cost control as part of ongoing platform stewardship, positioning consumption transparency as operational governance rather than advisory input.

Governed operational evolution: Tiger Analytics integrates Unity Catalog migration into ongoing managed services activity, linking governance enforcement to live environment management. Restricted dataset handling requires controlled access configuration within operations. The company's AI-driven triage workflows introduce prioritization discipline into support routines, while aligning optimization efforts with Databricks program initiatives.

ML and GenAI operationalization: Tiger analytics extends managed services beyond consulting and modernization to include ML and LLM lifecycle governance. It incorporates automation frameworks for migration validation, code quality checks, model monitoring and responsible AI guardrails within Databricks environments. This demonstrates capability in sustaining production-grade AI deployments and reducing operational complexity across evolving data estates.

Caution

Tiger Analytics can further formalize lifecycle automation by defining CI/CD promotion controls, rollback processes, structured observability and SLA-linked reliability metrics within its Databricks managed services model to strengthen platform maturity and multi-environment scalability.

ISG Provider Lens®

APRIL, 2026

REPORT: DATABRICKS ECOSYSTEM PARTNERS