

Magic Quadrant for AI Platforms for Data Science and Machine Learning

22 June 2026 - ID G00839857 - 51 min read

By Yogesh Bhatt, Afraz Jaffri, [and 1 more](#)

AI platforms for data science and machine learning provide software that supports end-to-end AI models and agent development and life cycle management using diverse data science and AI techniques. Leaders responsible for AI need to select platforms with awareness of native AI agents that accelerate model development and deployment.

Strategic Planning Assumptions

By 2027, 50% of data analysts will be retrained as data scientists, and data scientists will shift to AI engineers.

By 2027, organizations will implement small, task-specific AI models, with usage volume at least three times more than that of general-purpose large language models.

By 2027, there will be three times more open positions for AI engineers than for data scientists, primarily because the development of pretrained AI solutions is displacing the need for custom-made machine learning training.

By 2028, at least 15% of day-to-day work decisions will be made autonomously through agentic AI, up from 0% in 2024.

Market Definition/Description

Gartner defines AI platforms for data science and machine learning (DSML) as platforms that support end-to-end AI model and agent development and life cycle management using diverse data science and AI techniques. These platforms enable data preparation, model building, deployment and governance, and are delivered as fully managed cloud services or on-premises infrastructure for AI experts and business users.

AI platforms for DSML primarily target AI experts, including data scientists and AI engineers, providing a comprehensive suite of data science and AI techniques for augmented and automated decision making. These platforms support building models using data-intensive methods such as machine learning, as well as techniques such as simulation and optimization. They handle all types of data — tabular, image, video and text — for multimodal applications, such as computer vision, natural language processing (NLP) or composite AI, that combines optimization with machine learning.

Models can leverage both classic and modern AI techniques to automate business processes, such as credit scoring, churn prediction, predictive maintenance, recommendation, image classification, search and retrieval. They also enable the creation of AI agents that drive autonomous operations within AI-powered workflows and tasks across the DSML life cycle.

AI platforms for DSML foster collaboration and asset reuse across teams and departments while orchestrating workloads to manage large data volumes. They provide consistent, reproducible training and development environments that maintain lineage between data, code and models, improving productivity for AI experts. Additionally, low-code interfaces, natural language tools and AI assistance allow domain experts and business users to create predictive models through simple interactions.

These platforms support multiple Ops practices for deploying models and agents in production, including orchestration of both batch and real-time workloads, implementation of guardrails and AI governance for risk management, and continuous evaluation and monitoring of model and agent performance.

Mandatory Features

- Import or connect to a variety of structured, semistructured and unstructured data from data management systems, including databases, data warehouses, applications and content repositories located on-premises and in the cloud.

- Prepare data using data transformation tools and packages, including NLP and large language model (LLM) techniques, for data and metadata enrichment. This includes synthetic data generation, annotation and labelling.
- Provide a code-based development environment that supports software engineering practices and exploratory, interactive development.
- Build and evaluate models using a library of core statistical methods and generative AI techniques, including algorithms and processes.
- Enable collaboration and project management tools for multiple users and teams to work within the platform.
- Deploy, host and serve models in the platform for use in services and applications.
- Manage the life cycle of models and agents, including promotion, demotion, retraining and retirement.
- Offer administration and configuration management for user roles, permissions and resource allocation.

Optional Features

- Provide platform-generated recommendations for preparing, integrating and modeling data, along with automated machine learning model creation based on manually selected target predictions.
- Offer advanced interfaces that support complex modeling for simulation, optimization and deep learning use cases.
- Include custom software development kits (SDKs) for greater control and flexibility in code-based model development and integration with services and applications.
- Support methodologies and frameworks for handling structured and unstructured data sources, including text, images, video, audio and geospatial data.
- Deliver low-code interfaces and autoML functions for nonexpert roles, such as business users and domain experts.
- Enable postdeployment life cycle management to retrain, retire or adapt models based on data, feature and model drift.

- Provide tools and processes for operating data pipelines, ML models and AI agents at scale across diverse production environments.
- Offer functionality for working with generative AI models, including large language and reasoning models (LLMs and large reasoning models [LRMs]), through provisioning, selection, fine-tuning and prompt/output monitoring.
- Include techniques and tools that improve transparency, explainability and interpretability of models to clarify how outputs are generated.
- Support metadata management and cataloging through structured repositories of key assets, including data, code, features, models, logs and outputs.
- Enable specialized hardware support (e.g., GPUs) for training deep learning and generative AI models.
- Provide advanced data visualization for hypothesis testing, data validation, exploratory analysis and use-case identification.
- Integrate data lineage and provenance tracking to ensure data integrity and traceability throughout the workflow.
- Offer advanced hyperparameter tuning and optimization tools to improve model performance and efficiency.
- Support building LLM-based AI agents, including management of models, tools, memory, protocols and outputs.
- Enable retrieval-augmented generation (RAG) systems, including unstructured data preparation, vector search, embedding models and evaluation.

Magic Quadrant

Figure 1: Magic Quadrant for AI Platforms for Data Science and Machine Learning





Gartner

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud is a Challenger in this Magic Quadrant. Its Platform for AI (PAI) and Model Studio products are broadly focused on providing high-performance computing resources and infrastructure for data science teams to build, deploy, and manage ML models, GenAI models, and enterprise-grade agents. Its operations are mostly in China, Southeast Asia, and the Middle East, and it is expanding into Europe and the Americas, with clients comprising enterprises of various sizes across sectors. Its investments include a \$53 billion infrastructure commitment toward artificial superintelligence (ASI),

a strategic alliance with NVIDIA for physical AI, and advancing its agentic and Qwen open-source ecosystems.

Strengths

- **Vision:** Alibaba Cloud has a differentiated outlook on the advancement of AI models and agents, combining physical AI, simulation, and self-learning. It is investing significant capital to achieve this vision, enabling customers to pioneer advanced, future-ready enterprise AI applications.
- **Research and innovation:** Alibaba Cloud continues to release new versions of Qwen, its series of open-weight models, that match or exceed commercial benchmarks. It was also awarded best paper at the NeurIPS 2025 conference. These developments show how Alibaba Cloud provides customers with access to the latest technologies for scaling their AI initiatives.
- **Model tuning:** Alibaba Cloud's products offer a wide range of model-tuning features for training, testing, and inference optimization to create domain-specific small, large, and multimodal language models. These features allow Alibaba Cloud customers to efficiently deploy highly customized models tailored to their unique business use cases.

Cautions

- **Learning curve:** The Alibaba Cloud ecosystem has a steep learning curve, especially for nontechnical users, to derive full value. Learning how to use the platform requires a time- and resource-intensive commitment that may delay an organization's AI return on investment.
- **Multicloud support:** Options for using Alibaba Cloud as part of a multicloud strategy are limited because the provider lacks intercloud partnerships with other major providers for data access and reuse. Customers relying on distributed cloud ecosystems may face significant friction and increased costs when integrating their data.
- **Geographic presence:** Large multinational organizations can experience challenges with Alibaba Cloud when maintaining enterprise solutions in conformance with sovereign AI mandates, despite the fact that the vendor has data and compute centers globally and offers Apsara Stack for on-premises deployments. Customers in highly regulated regions may struggle to satisfy strict local data and compliance rules.

Amazon Web Services

Amazon Web Services (AWS) is a Leader in this Magic Quadrant. Its Amazon SageMaker AI and Amazon Bedrock products are mainly focused on enabling data science and machine learning (DSML) model building, developing GenAI applications, and scaling production-ready AI agents. Its operations are geographically diversified, and its clients comprise organizations of all sizes and sectors. Recent investments include the AWS Trainium3 chip for better token economics in next-generation agentic, reasoning, and video applications. It also became the exclusive third-party cloud distributor for OpenAI Frontier to manage AI agent teams and released the Amazon Nova 2 Lite reasoning model for cost-efficient everyday tasks.

Strengths

- **Partner ecosystem:** AWS' extensive provider partner ecosystem enables best-of-breed products to be easily integrated into SageMaker AI and Amazon Bedrock. This allows AWS customers to achieve greater innovation and flexibility in their product innovation and usage.
- **AgentOps:** AWS has bundled its key AI agent components into the Amazon Bedrock AgentCore framework for managing identity, observability, evaluation, memory, tools, and runtime. This helps AWS customers efficiently build, manage, and deploy scalable and reliable agentic systems.
- **Scalability:** AWS continues to invest in AI-optimized infrastructure for the most demanding training and inference workloads, including Trainium3 and SageMaker HyperPod, and partnerships with providers such as Cerebras. This ensures customers have the high-performing computing required to scale their enterprise AI initiatives.

Cautions

- **Product offering:** The separation between Amazon Bedrock and SageMaker AI remains, even though AWS has made significant progress in bundling SageMaker components into a unified offering. This separation between products causes friction for cross-disciplinary delivery teams, who may struggle to seamlessly integrate their workflows.
- **AI governance:** Features such as Amazon Bedrock Guardrails, data quality, and drift monitoring are available, but the platform's usage for overarching risk and policy

compliance is limited. As a result, AWS customers in highly regulated industries may need to invest in third-party governance tools.

- **Cost management:** AWS customer feedback frequently highlights challenges with managing costs and navigating pricing complexity, particularly for large-scale workloads or long-running instances. Although AWS is introducing additional budgetary tools, customers must carefully assess the accuracy of these tools to ensure effective cost management and avoid unexpected expenses.

Cloudera

Cloudera is a Visionary in this Magic Quadrant. Its Cloudera AI and Cloudera Unified Data Fabric products are broadly focused on providing a unified data platform to build, deploy, and govern mission-critical data, models, and agents across hybrid clouds and data centers. Its operations are geographically diversified, and its clients primarily comprise organizations across banking, telecommunications, auto manufacturing, and government. Future investments include advancing Cloudera Agent Studio for multiagent workflows, scaling its on-premises and virtual private cloud AI Inference service via NVIDIA and Dell partnerships, and integrating Taikun as its core Kubernetes platform, building also on its acquisitions of Verta AI Ops and Octopai (for data lineage).

Strengths

- **Vision:** Cloudera has converged its data and AI environments to utilize resources from on-premises and the cloud, resulting in its Unified Data Fabric offering. This permits Cloudera customers to cohesively integrate and scale their data and AI operations across large-enterprise workloads.
- **Geographic presence:** Cloudera has expanded operations in multiple regions, including the Middle East and Asia, and has developed numerous partnerships with local system integrators. This enables Cloudera customers to receive dedicated focus and support for their specific regional requirements.
- **Performance:** The vendor's acquisition of Taikun boosts scalability and efficiency for managing data and AI Kubernetes workloads across cloud environments and data centers. These enhanced capabilities help Cloudera users maintain high performance and reliability as the organization's complex data infrastructure grows.

Cautions

- **Retrieval-augmented generation (RAG):** Cloudera has limited native features to support deep research and other use cases, although the data platform offers vector store integrations. Consequently, Cloudera customers may need to invest in third-party tools to support advanced GenAI applications for search and retrieval.
- **Complexity:** Cloudera's products are designed as enterprisewide solutions that require implementation and customization by IT teams. Customers looking for targeted point solutions or those starting with smaller AI projects may experience higher initial investments.
- **Solution delivery:** Cloudera's internal catalog of industry-specific use cases is relatively limited in volume compared to competitors, so teams looking to build end-to-end solutions may need to hire dedicated professional services. This limitation can force Cloudera customers to rely more heavily on external talent for complex tailored implementations and platform maintenance.

Databricks

Databricks is a Leader in this Magic Quadrant. Its Data Platform product is broadly focused on offering integrated lakehouse, lakebase, governance, analytics, model training, and serving capabilities, including streamlined building of GenAI models and AI agents with Agent Bricks. Its operations are geographically diversified, and its clients tend to be midsize and large enterprises across sectors. Recent acquisitions include Neon for serverless Postgres, Tecton for real-time data serving, Fennel AI for feature management, and Quotient AI for AI agent evaluation. Future investments focus on expanding Agent Bricks and Unity Catalog to provide the control plane to build, monitor, and centrally govern enterprise agents.

Strengths

- **Innovation:** Databricks continues to add new features and functionality to its platform at an accelerated rate, focusing on data and unified governance as a foundation for AI agents. This helps Databricks customers overcome common adoption challenges and successfully deploy enterprise agentic workflows.
- **Operations:** The company exhibits strong growth and leadership stability alongside an increasing headcount across key engineering, delivery, and support roles. This level of organization ensures that Databricks customers can confidently make multiyear strategic investments in the platform.

- **Community:** The platform maintains strong adoption among data and AI practitioners, evidenced by a growing ecosystem of more than 69,000 tech-certified individuals and engagement through open-source initiatives. This allows Databricks customers to readily source the specialized talent required to work on the platform.

Cautions

- **Pace of innovation:** The volume of Databrick's updates and new features can overwhelm teams and cause a burden for administrators. Additionally, some of its documentation is incomplete, which can require heavy experimentation before productionizing. This may cause Databricks users to experience delayed timelines for enterprise AI deployments.
- **Composite AI:** While Databricks supports standard ML and GenAI within its platform, combining them relies on code-centric approaches rather than prebuilt workflows. Furthermore, techniques like simulation utilize the platform's extensive open-source integrations instead of native proprietary features. This approach requires customers to build and tailor their composite AI integrations.
- **Hybrid workloads:** Databricks does not offer a hybrid solution, as it only supports workloads running in cloud environments. Due to this, Databricks is not suitable for buyers who need to run workloads on-premises.

Dataiku

Dataiku is a Leader in this Magic Quadrant. Its Platform for AI Success is mainly focused on combining low-code development and AI-assisted tools with code-first interfaces to enable collaborative analytics, DSML, governed agent orchestration, and reliable composite AI applications. Its operations are geographically diversified across the Americas, Europe, and Asia/Pacific, and its clients tend to be midsize and large enterprises across all sectors. Future investments focus on advancing Dataiku Agent Management, Dataiku Cobuild, and Dataiku Reasoning Systems to shift from isolated agents to interconnected decision workflows.

Strengths

- **Vision:** Dataiku's development of new reasoning systems represents an advanced understanding of enterprise pain points and a convergence of DSML and decision

intelligence. With these advancements, Dataiku customers are well-positioned to successfully integrate agentic processes into their business operations.

- **Flexibility:** Dataiku's platform acts as an orchestration layer that manages data flows and agent actions across disparate data sources and platforms. As a result, Dataiku customers benefit from architectural flexibility and avoid lock-in to a single end-to-end stack.
- **Risk and governance:** The platform provides integrated AI governance that supports various personas across AI leadership, covering AI audits, portfolio supervision, and use-case-specific controls. This ensures that Dataiku customers can effectively manage risk and compliance across their enterprise AI initiatives.

Cautions

- **Competition:** Like all independent vendors, Dataiku faces tough competition from data management and business application providers that are expanding their AI capabilities to own more of the technology stack. Dataiku's platform operates independently, which may not be suitable for enterprises seeking a single-vendor architecture.
- **Cost:** Dataiku customer feedback often cites high overall costs and licensing models as limiting factors when it comes to contract renewal decisions. To achieve greater ROI with the platform, Dataiku customers must ensure that they fully utilize advanced capabilities, rather than relying solely on the product as a basic low-code tool.
- **Differentiation:** In an expanding market for agent building tools, particularly low-code entrants, Dataiku could do more to help buyers understand its platform's unique value. For organizations starting with basic AI workflows, lower-cost alternatives may present an easier path to adoption than Dataiku.

DataRobot

DataRobot is a Leader in this Magic Quadrant. Its Agent Workforce Platform, co-engineered with NVIDIA, is broadly focused on delivering, managing, and monitoring predictive AI, agentic AI and GenAI solutions, with composable ML for advanced data scientists, and MLOps and LLMOps. It specializes in high-complexity agentic workflows across hybrid, multicloud, on-premises, and air-gapped environments. DataRobot's operations are geographically diversified, and its clients tend to be large global

enterprise leaders. Future investments focus on helping organizations meet sovereign AI mandates and automating core business processes via agentic transformation. It also partners with NVIDIA, Dell, and Nebius to offer a preinstalled agent workforce platform with full build and runtime capabilities outside of public clouds.

Strengths

- **Deployment options:** The Agent Workforce Platform supports AI workloads across cloud, hybrid, on-premises, and air-gapped environments. Moreover, the integration of Nebius presents an alternative for hyperscalers. Consequently, DataRobot customers gain extensive architectural flexibility and avoid provider lock-in.
- **Customer experience:** DataRobot customers consistently provide positive feedback regarding the high level of service and support that the vendor offers throughout the platform adoption phases. This approach substantially reduces implementation risks and increases investment confidence for DataRobot customers.
- **Partnerships:** DataRobot's strategic partnership with NVIDIA integrates the platform into the Enterprise AI Factory architecture, granting users the ability to leverage NVIDIA's tools for agent management. Ultimately, DataRobot buyers benefit from simplified integration and evaluation workflows.

Cautions

- **Community:** DataRobot provides limited avenues for users to share best practices and feedback on product features compared to other providers. Ultimately, this extends the learning curve and hinders the ability of DataRobot customers to self-resolve common implementation issues.
- **Cost:** Cost remains a concern for DataRobot buyers, particularly at usage extremes for smaller organizations or those deploying the platform widely. Since the effects of the new pricing model are not yet fully understood, DataRobot buyers must carefully forecast their specific usage thresholds to avoid unexpected costs.
- **Ecosystem:** DataRobot's platform requires sustained integration support from provider partners in order for it to act as a central agent platform layer across enterprise applications. Without additional strategic partnerships like SAP, DataRobot customers risk feature duplication and integration setbacks.

Domino Data Lab is a Visionary in this Magic Quadrant. Its Enterprise AI Platform product, comprising Model Factory, App Hub, and Governance Center, broadly focuses on delivering an open, adaptable enterprise AI platform to build, deploy, and govern AI, GenAI, and agentic systems across any environment. Its operations are geographically diversified, and its clients tend to be global enterprises in highly regulated industries like life sciences, financial services, and the public sector. Future investments focus on facilitating the operation of AI agents with integrated governance, reproducibility, and security for the full spectrum of AI use cases, alongside improving the developer experience by integrating leading AI coding agents into DSML workflows.

Strengths

- **Observability:** In addition to strong governance capabilities, Domino's platform provides end-to-end traceability as a system of record across AI development, integrating with open-source and commercial frameworks to deliver comprehensive monitoring, evaluation, and governance for model and agent use cases. This allows Domino customers to effectively manage risk and maintain visibility across their AI portfolios.
- **Market understanding:** The Enterprise AI Platform focuses on the needs of data scientists and AI engineers across the development life cycle, especially for MLOps, compute management, and agentic tooling. Domino's platform reduces the need for specialist operations personnel, enabling delivery teams to take end-to-end ownership of development.
- **Agents for data science:** Domino's platform offers agents integrated into the DSML life cycle, and it enables external agents to use the platform as a tool via Model Context Protocol (MCP) integration. These agents accelerate time to value for data scientists and provide architectural flexibility when building multiagent systems that require access to internal Domino assets.

Cautions

- **Complexity:** Domino customers frequently operate in complex regulated environments and cite integration and deployment of the platform as common concerns. While the addition of embedded AI functionality is intended to reduce complexity for IT teams over time, prospective Domino buyers must plan for extended implementation timelines and allocate sufficient internal resources.

- **Industry strategy:** The provider specializes in key highly regulated industry verticals that have requirements aligning with its product's support for highly complex and regulated environments. Organizations operating outside of these specific sectors may find that Domino is less aligned to facilitate their general-purpose enterprise needs.
- **Partnerships:** The provider maintains alliances with technology providers and global system integrators, but its ecosystem is more curated and is not as strategically supported as other market competitors. Without a robust partner network, Domino customers face a limited availability of resources for platform deployments and solution development.

Google

Google is a Leader in this Magic Quadrant. Its Vertex AI (incorporated into Gemini Enterprise Agent Platform as of April 2026) product is an integrated offering for DSML, GenAI, and AI agent use cases, mainly focused on providing enterprise teams with a unified, secure foundation for data science and MLOps. Google's operations are geographically diversified (excluding China), and its customers are enterprises of all sizes across industries. Future investments focus on expanding capabilities across the AI stack, including next-generation AI accelerators, advancing its family of AI models (including Gemini) and enhancing Vertex AI Agent Builder. Google also drives open agentic standards via the A2A protocol and plans between \$175 billion and \$185 billion in 2026 capital expenditure, mainly for AI infrastructure and full stack strategy.

Strengths

- **AI offering:** Vertex AI does not depend on any other provider for any components across the entire AI stack, from infrastructure to models and applications. This independence presents a compelling offering for enterprises seeking a strategic partner for multiyear planning or looking to maximize existing Google investments.
- **Developer experience:** Google has released numerous features to aid AI engineers during the development process, including purpose-built data agents, data science agents, and agent skills. By simplifying the creation of data engineering, data science, and AI products, Google allows organizations to accelerate their time to market.
- **Agent ecosystem:** The provider maintains a focus on enabling agent interoperability across platforms by pioneering the A2A protocol and adding more than 150 partners

to coordinate workflows and communication. Google's ecosystem provides customers with the flexibility to build autonomous workflows that touch multiple systems.

Cautions

- **Composite AI:** Using traditional ML and diverse AI techniques in Vertex AI can be more challenging than in other platforms, as the platform gives prominence to LLM-based agent building and GenAI application development. Google customers may find it difficult to deliver complex composite AI systems with Vertex AI.
- **Cost:** Customer feedback frequently cites unpredictable costs as an issue with using Google's services, and its varied tier-based pricing makes forecasting a challenge. While Google offers a comprehensive suite of cost management capabilities, including FinOps, some customers may be limited when managing the costs of workloads outside of Google Cloud.
- **Coherence:** Google's rapid expansion of its portfolio has resulted in overlapping capabilities between tools, such as Vertex AI Search, Vertex AI Agent Engine, Google Antigravity, Colab Enterprise, Google Agentspace, and the overarching Vertex AI Agent Builder. This sprawling portfolio can be confusing to prospective Google customers seeking to build tailored enterprise solutions.

H2O.ai

H2O.ai is a Visionary in this Magic Quadrant. Its products include h2oGPTe for agentic use cases, H2O Driverless AI for predictive analytics, and H2O AI Super Agent for complex scenarios. It broadly focuses on delivering sovereign, on-premises, and air-gapped AI to build, fine-tune, and automate domain-specific agentic workflows. Its operations are geographically diversified across APAC, EMEA, and LATAM, and its clients tend to be large enterprises in financial services, telecommunications, healthcare, and the public sector. Future investments focus on evolving its AI Super Agent platform for industry use cases, increasing low-code features in Agents Builder, and updating H2O-3, its open-source ML framework.

Strengths

- **Trustworthy AI:** H2O.ai's platform is used extensively in regulated industries across complex use cases, featuring transparent security controls, compliance measures,

and achievement of FedRAMP High authorization. This focus ensures H2O.ai customers can safely operate within highly sensitive environments.

- **Customer collaboration:** The provider's forward-deployed engineers (FDE), Kaggle Grandmasters, and extensive community reduce the barriers for platform adoption. This allows H2O.ai users to achieve ongoing transformation through AI.
- **User experience:** H2O.ai's platform caters to business roles building agentic solutions as well as data scientists and AI experts. With a low learning curve and advanced features for model fine-tuning, RAG, and agent observability, H2O.ai enables diverse users to efficiently construct complex enterprise workflows.

Cautions

- **Viability:** H2O.ai profitably competes against much larger providers with full stack offerings as well as those with a highly specific focus. The continued R&D investment required to keep pace with its competition may be challenging for H2O.ai to sustain. H2O.ai customers may face potential risks regarding the platform's long-term evolution.
- **Visibility:** Gartner sees a low volume of customer interest in H2O.ai's platform compared to other evaluated providers. This lack of broad market visibility makes it harder for prospective buyers to confidently identify whether H2O.ai is the best fit for their specific enterprise use case.
- **Industry solutions:** The platform primarily focuses on highly regulated sectors and continues to expand its playbook coverage across all regulated verticals; but, for some of these industries, it does not yet offer specific solution playbooks and architectures. H2O.ai users may experience an extended timeline to achieve full-scale adoption.

IBM

IBM is a Leader in this Magic Quadrant. Its watsonx portfolio is an integrated offering supporting data management, unified governance, and GenAI. It is focused on enabling organizations to build, deploy, and orchestrate scalable multiagent systems across complex hybrid cloud architectures. Its operations are geographically diversified via a large global partner network, and its customers comprise large enterprises across sectors. Its 2026 go-to-market strategy shifted from product-led to outcome-driven,

targeting specific roles within customer accounts. Future investments focus on developing prebuilt agents and tools, expanding AgentOps capabilities, and optimizing the functionality of AI Gateway.

Strengths

- **Governance:** IBM offers a comprehensive suite of AI governance capabilities through watsonx.governance. This offering not only governs models and agents built with watsonx but also monitors and reports on third-party services, and supports sovereign AI requirements, reducing the need for separate governance tools across the enterprise.
- **Foundation models:** The watsonx.ai platform features numerous purpose-built models for enterprise applications, including embedding, multimodal and domain-specific use cases, and also supports third-party and open-source models. By offering options that are often easier and more cost-effective than frontier models, IBM allows buyers to efficiently scale AI initiatives.
- **RAG:** IBM's system provides multiple types of RAG techniques, including GraphRAG and Agentic RAG, making it easier to address deep research use cases. By combining structured and unstructured data with intelligent agents, IBM empowers enterprises to build highly effective search and retrieval solutions.

Cautions

- **Context enrichment:** IBM's agent building and orchestration products have limited capabilities for the inclusion of a context layer spanning semantics and knowledge graphs. As a result, IBM customers building multiagent systems may need to source these capabilities from other providers or partners.
- **Fragmentation:** The watsonx portfolio consists of multiple products including watsonx.data, watsonx Orchestrate, watsonx.ai and watsonx.governance. Although this provides flexibility for single product use, IBM customers utilizing the full suite may experience initial confusion about which product provides which capabilities.
- **Data science agents:** The provider offers its own coding agent, Bob, that automates software tasks, but it lacks the dedicated agents for data modeling, machine learning, and operationalization found in competing platforms. To increase data science productivity with AI agents, IBM customers may need to use external providers.

MathWorks

MathWorks is a Niche Player in this Magic Quadrant. Its MATLAB and Simulink products comprise an integrated offering that is broadly focused on accelerating Model-Based Design, system simulation, and AI deployment for complex engineered systems. Its operations are geographically diversified, and its customers tend to be within specialized sectors, such as the automotive, aerospace, and robotics industries. Future investments focus on continuing to extend its AI capabilities across both platforms, with updates to MATLAB Copilot, embedded and edge AI, and AI for safety-critical systems.

MathWorks did not respond to requests for supplemental information or to review the draft contents of this document. Gartner's analysis is therefore based on other credible sources.

Strengths

- **Engineering AI and optimization:** For users looking to train and build their own AI agents that interact with physical environments, MathWorks offers highly specialized options like the Reinforcement Learning Toolbox. This comprehensive suite of tools helps MathWorks customers successfully develop robotics and autonomous systems.
- **Simulation:** MATLAB excels at managing specialized engineering data and integrates deeply with Simulink to systematically test AI models within simulated physical environments. By facilitating accurate evaluation, MathWorks allows customers to reliably test performance without committing to physical investment.
- **Edge AI:** MathWorks' products facilitate rapid prototyping through intuitive low-code apps and allow for efficient conversion of trained models into deployable code for edge devices and embedded hardware. This allows MathWorks customers with non-AI backgrounds to successfully produce embedded systems with AI capabilities.

Cautions

- **Orchestration:** The provider does not focus on enabling LLM-based multiagent systems to be orchestrated by its platform. MathWorks customers building these architectures will need to rely on external platforms to provide this functionality.
- **Application development:** MathWorks is primarily focused on supporting highly specialized engineering use cases intended for use in embedded systems. To

successfully build applications and dashboards that utilize trained models, MathWorks customers will need to depend on external integrations.

- **Ecosystem:** The provider maintains partnerships with cloud computing providers but has not established a broader ecosystem of data management or DSML providers. Without these native connections, MathWorks customers may need to create custom integrations to manage workloads across different providers.

Microsoft

Microsoft is a Leader in this Magic Quadrant. Its Azure Machine Learning (Azure ML) product is part of Microsoft Foundry platform. It provides data scientists, AI engineers, and developers with an AI application and agent factory for model development, AutoML, and MLOps, and building action-oriented, context-aware agents at scale. Its operations are geographically diversified, and its customers operate across various sizes and sectors. Future investments focus on accelerating time to value through the expansion of productivity-improving agents, increasing interoperability with tools and frameworks of choice, and building scalability through purpose-built AI infrastructure.

Strengths

- **GenAI models:** Microsoft offers an integrated ecosystem, providing access to more than 11,000 foundation models from diverse partners, including OpenAI, Anthropic, and Mistral. Alongside more than 1,400 prebuilt connectors and Model Context Protocol (MCP) tools, Microsoft empowers customers to connect third-party enterprise applications.
- **Context layer:** Fabric IQ is positioned as providing the necessary context for AI and agent-based applications to operate within a complex enterprise environment. By equipping customers with an integrated knowledge layer, Microsoft facilitates the production of reliable systems.
- **Personas coverage:** Azure ML caters to a broad range of user personas across development journeys, including data, DSML, AI engineering, and governance. By facilitating seamless collaboration across these groups, Microsoft aids customers in successfully reducing silos between teams.

Cautions

- **Stability:** Many of Microsoft's new AI features remain in preview mode for extended periods of time. This elongated timeline may cause Microsoft customers to hesitate when utilizing new capabilities in production-ready deployments.
- **Contract negotiation:** Customer reports on Microsoft often cite negative experiences regarding contract evaluation and negotiation in comparison to interactions with competing providers. Microsoft customers may need to apply extra effort to arrive at the best deal for their specific requirements.
- **Product strategy:** Microsoft promotes Foundry as its default platform for enterprise AI agents, while continuing to promote Azure ML's roadmap for traditional MLOps. This dual approach creates overlapping messaging, which risks architectural and branding confusion among prospective Microsoft customers.

Posit

Posit is a Niche Player in this Magic Quadrant. Its Posit Connect, Positron IDE, Posit Workbench and Posit Package Manager products form an integrated offering. It is broadly focused on realizing value from the full DSML workflow using a code-first approach, with support for both R and Python, and a large variety of open-source libraries to ensure verifiable, auditable, and production-grade models and apps. Its operations are geographically diversified across the Americas, EMEA, and APAC, and its customers tend to be large enterprises across all sectors. Posit's focus is on promoting Positron IDE as a neutral control plane with governed connections to multiple data platforms and providing support for industry-specific solutions to accelerate time to value.

Strengths

- **Corporate strategy:** The provider is a public benefit corporation that focuses on serving academia and not-for-profit organizations, contributing to research and open-source communities beyond its primary business serving commercial enterprise customers. This unique strategy enables Posit customers to align their technology investments with enterprise social responsibility policies.
- **Core data science:** Posit focuses on analytical data science and generating insights for decision making with AI in high-stakes use cases. By providing intuitive analytical tools and AI-driven insights, Posit enables customers to simplify the data-to-action decision-making process so they can realize value faster.

- **Industry focus:** The provider has established a strong presence in highly regulated sectors such as financial services and life sciences. Posit customers in these industries have access to a set of mature platform capabilities for compliance-grade AI deployment and clinical innovation.

Cautions

- **AI engineering:** The provider focuses on insight generation, orchestrates compute and relies on partnerships with other providers to serve use cases with extensive data, compute, or infrastructure requirements. By introducing integration complexity for GenAI and agentic applications, this approach may lead Posit customers to experience delays.
- **Personas:** The provider is mainly suitable for organizations with mature data science domain experts who possess a coding or statistical background, while adding capabilities to lower the entry barrier. Due to this highly specialized focus, Posit customers who are just starting with data science initiatives may experience a steep learning curve that hinders overall adoption
- **Collaboration:** Due to the code-first nature of Posit, seamless collaboration between multiple personas and comprehensive workflow support are limited compared to competing offerings. To adequately provide this critical functionality, Posit customers may need to invest in external project management tools.

Red Hat

Red Hat is a Visionary in this Magic Quadrant. Its OpenShift and AI Enterprise products provide a scalable, open, and secure platform for deploying, serving, and managing AI models and agentic workflows across hybrid cloud and on-premises settings. Its operations are geographically diversified, and its customers are typically global enterprises across all industry sectors. The provider has made a number of recent acquisitions to develop its portfolio, including Neural Magic for model serving, Jounce for observability, and Chatterbox Labs for safety and risk assessment against adversarial attacks.

Strengths

- **Flexibility:** Red Hat provides a range of options for enterprises looking to use their own infrastructure for sensitive use cases, from physical bare metal to edge

deployments. This allows customers to securely manage their environments while adhering to strict compliance rules.

- **Scalability:** The acquisition of Neural Magic provides direct access to vLLM and Llm-d, enabling Red Hat to provide distributed, optimized, and cost-effective GenAI inference capabilities across hybrid cloud environments. This ensures customers can efficiently maintain high-performing operations as their complex applications grow.
- **Innovation:** Through co-engineered partnerships with major hardware and cloud providers like NVIDIA, AMD, and AWS, Red Hat delivers validated, ready-to-use, open AI Factory solutions. These features prevent provider lock-in and standardize enterprise support for agentic AI, accelerating time to value.

Cautions

- **Complexity:** To effectively set up, operate and maintain Red Hat's products, customers will likely need dedicated infrastructure teams. Without sufficient engineering resources, potential Red Hat customers may experience extended implementation timelines and struggle to properly use the platform.
- **Composite AI:** Red Hat's products are more recently less focused on traditional ML, as it is centered on GenAI and LLM-based agents. This constraint presents a challenge in delivering composite AI, potentially forcing Red Hat customers to defer to external platforms to build multimodal applications.
- **Partner ecosystem:** The integration with other platforms that complement the system's strengths is still in progress. Since there is less visibility from system integrators and consultancies, Red Hat customers may face limited availability of specialized resources for complex deployments and solution development.

SAS

SAS is a Visionary in this Magic Quadrant. Its Viya platform involves the entire data and AI life cycle, providing low-code and pro-code interfaces for governed, enterprise-grade AI development. Its operations are geographically diversified across the Americas, EMEA, and APAC, and its customers are mainly large enterprises in highly regulated sectors like banking and government. SAS' future investments focus on enabling users to build intelligent agents through composite AI techniques, providing advanced agents for data

science tasks, and solidifying observability with centralized agent orchestration and governance.

Strengths

- **Vision:** SAS uses composite AI agents and orchestration as a primary means to achieve intelligent decision making in mission-critical scenarios. This approach differentiates SAS from competing providers who are focused solely on GenAI models, and it helps SAS customers build highly reliable and accurate AI systems.
- **Composable architecture:** The provider continues to enhance its composable architecture by offering lightweight, independent entry points, such as SAS Viya Workbench and SAS Data Maker, through major cloud marketplaces. By utilizing these stand-alone modules, SAS customers gain flexible, low-cost environments for AI development.
- **Service and support:** SAS receives positive feedback from existing customers regarding its level of service, support, and overall buying experience. Customers adopting these offerings have multiple distinct avenues to receive support. This comprehensive coverage significantly reduces deployment friction, aiding SAS customers with seamless integration and adoption.

Cautions

- **LLM management:** SAS Viya lacks a native model catalog to centrally manage and connect to LLMs from external providers or on local instances. Connecting external or locally hosted LLMs is handled outside the platform, so organizations using multiple providers may want to plan for some additional coordination
- **Data engineering:** SAS' Viya platform centers processing on its proprietary Cloud Analytic Services (CAS) engine, alongside native SAS Compute, offering pushdown options for select data stores. While the platform supports embedding R and Python as scripting languages, SAS customers who use open-source frameworks may struggle to properly integrate the platform into their existing pipelines.
- **Context layer:** While agents can be built within SAS Viya, the system offers limited exposure to context-ready assets, such as semantic definitions, metadata, and knowledge graphs. This constraint makes it substantially more complex to orchestrate interoperable multiagent systems, which may push customers to rely on external tools.

Siemens (Altair)

Siemens (Altair) is a Visionary in this Magic Quadrant. Its AI Fabric platform (acquired from Altair in March 2025) consists of Rapidminer portfolio and Mendix. These products are mainly focused on building AI-driven automation, agentic workflows, and composite applications using enterprise knowledge graphs. Its operations are geographically diversified, and its customers are large enterprises across industrial, automotive, healthcare, and financial sectors. Siemens' 2026 strategic focus is to unify Rapidminer with Mendix to operationalize agentic intelligence while expanding ecosystem partnerships with Databricks and Snowflake.

Strengths

- **Manufacturing AI:** Siemens possesses deep domain expertise in deploying AI across highly complex industrial scenarios, utilizing the newly acquired Altair portfolio to deliver end-to-end hardware and software architecture. By leveraging this extensive manufacturing knowledge, Siemens enables customers to confidently implement specialized solutions that optimize their production processes.
- **R&D:** Siemens maintains a world-class research unit featuring €6.6 billion in R&D expenditures, more than 50,000 employees, and 2,600 patents to produce unique outputs like the Industrial Foundation Model. This extensive commitment to innovation ensures Siemens customers can continuously leverage cutting-edge, domain-specific technologies to speed up their complex deployments.
- **Full stack offering:** Siemens ongoing integration of the platform's AI fabric and machine learning tools with Mendix will deliver a comprehensive application development suite. By uniting diverse engineering teams within a single ecosystem, this end-to-end architecture will allow Siemens customers to seamlessly transition from raw data ingestion to automated execution.

Cautions

- **Coherence:** While the functionality of Siemens' products across Rapidminer, Mendix and the rest of its portfolio is being actively unified under the AI Fabric architecture, customers may need time to fully adopt the integrated stack.
- **Industry focus:** Siemens is best known for manufacturing and engineering sectors, though Rapidminer has its origins in adjacent verticals such as banking and financial

services. Expanding the breadth of industry-specific support beyond its core sectors remains an area of opportunity.

- **Deep research:** While the platform's knowledge graph component within Rapidminer provides MCP integration and automated ontology creation for agentic exploration, more advanced RAG and deep research use cases may still require additional configuration effort.

Snowflake

Snowflake is a Visionary in this Magic Quadrant. Its Cortex AI platform features Snowflake CoWork (formerly Snowflake Intelligence) and Snowflake CoCo (formerly Cortex Code) products that broadly provide a unified cloud-based platform for DSML practitioners to support data engineering, analytical workflows, and developing secure agentic AI applications. Its operations are geographically diversified, serving midsize-to-large enterprises across all industries. Snowflake's recent acquisitions of Observe, for AI-powered observability, and TensorStax accelerate agentic data engineering. The provider's 2026 roadmap focuses on building an AI Data Cloud and enterprise agent architecture that provides context, decision intelligence, and governance for systems that can reason across business processes.

Strengths

- **Vision:** Snowflake maintains a strong focus on enabling AI agents and overcoming inherent GenAI limitations by providing context and decision intelligence. By using these foundations, Snowflake customers can use the platform for advanced use cases and multiagent systems.
- **Data agents:** The platform's CoCo operates as a truly agentic system, bringing advanced features typically seen in software engineering directly to data engineering and data science workflows. By maximizing this functionality, Snowflake customers benefit from notable improvements in automation and overall team productivity.
- **Operations:** Snowflake is actively scaling its support functions to deliver deep engineering, research, and industry experience through a dedicated forward-deployed engineer program. Engaging directly with these specialized experts allows Snowflake customers to benefit from decreased time to value and knowledge sharing.

Cautions

- **Composite AI:** The provider supports machine learning and GenAI features but relies on open-source libraries to implement alternative AI techniques. Snowflake customers using the product for building composite AI systems may need to use other vendors' products.
- **Hybrid environments:** Snowflake offers connectors to enable access of data across multicloud and on-premises environments, but usage of features such as Cortex AI is only available within Snowflake's managed infrastructure.
- **Complexity:** Customer feedback to Gartner frequently cites the complexity of the platform as being a concern. Prospective buyers need to plan for training and time to ramp up for effective use of the platform.

Teradata

Teradata is a Visionary in this Magic Quadrant. Its AI Studio with Teradata Enterprise AgentStack capabilities deliver a flexible, scalable, and autonomous AI and knowledge platform for driving autonomous agent workflows. Its operations serve large enterprises across all sectors and geographies. The provider's strategic focus centers on providing capabilities for the end-to-end agent life cycle across hybrid data environments, including a context engine and AgentOps capabilities.

Strengths

- **Unified platform:** Teradata offers AI Studio to integrate analytics, vector search, model management, and agentic automation in support of enterprise-scale AI use cases. Customers can build end-to-end applications without the need for other solutions.
- **Vision:** Teradata Enterprise AgentStack features AgentEngine, AgentBuilder, and AgentOps functionality that enables customers to quickly build, deploy, and manage agents governed by corporate policy boundaries. This additional context layer provides differentiation from competing solutions and improves the viability of agentic systems.
- **Open approach:** Teradata embraces an open approach by supporting deep integration with popular open-source Python ML frameworks and open data standards for customers that wish to build their own solutions. This enables customers to use their tools of choice while working directly with governed enterprise data.

Cautions

- **Investment scale:** The provider does not have the same scale of funding as some larger competitors, which requires a more focused approach to AI innovation through partnerships, acquisitions, or internal development. Customers should monitor Teradata's product roadmap and acquisition strategy to ensure it aligns with their needs.
- **Foundation models:** Teradata has chosen an open-model strategy rather than delivering a native-hosted offering within its own cloud environment. It provides access to foundational model provider services through APIs. Customers will need to set up and manage access to these provider services.
- **Cost Management:** Customer feedback cites that the provider's multilayered pricing model, which is designed to support increasing usage, is difficult to predict. Teradata has introduced unified unit-based pricing to simplify consumption. Customers should ensure they are using Teradata's latest pricing model and cost-management approach.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

Added

- Posit
- Red Hat
- Siemens (replacing Altair postacquisition)
- Teradata

Dropped

- Alteryx

Inclusion and Exclusion Criteria

To qualify for inclusion, vendors need to provide:

- A go-to-market strategy for their AI platform for DSML that focuses on the persona of a professional data scientist, machine learning engineer, or AI engineer.
- A DSML platform that is recommended for use in predominantly code-centric data science teams.
- At least 10 paying customers (logos) as of 1 January 2026 for their DSML platform in each of two of the following regions:
 - North America
 - South America
 - Europe, the Middle East and Africa
 - Asia/Pacific

Evaluation Criteria

Gartner evaluates a vendor's Ability to Execute in the AI platforms for data science and machine learning market by applying criteria that assess its capability to address current market trends. This involves demonstrating a high-quality track record of maintaining visibility and brand awareness in the AI space and operational practices that can attract and sustain high levels of customer demand and engagement.

Ability to Execute

Product or Service: This criterion covers the assessment of the vendor's capabilities to deliver features and functionality within its platform for core data science and ML-engineering-based use cases.

Overall Viability: This criterion covers the assessment of a vendor's key financial growth metrics and diversity of its customer base.

Sales Execution/Pricing: This criterion covers the assessment of customers’ willingness to make strategic investments in a vendor and the vendor’s ability to support client interest or demand.

Market Responsiveness and Track Record: This criterion covers the assessment of the quality of a vendor’s release cycles for its product, how it responds to customer needs, and the ability to prioritize feature requests.

Marketing Execution: This criterion covers the assessment of a vendor’s market visibility, brand awareness, and community development.

Customer Experience: This criterion covers the assessment of a vendor’s quality of technical support, level of customer retention, and customer advocacy.

Operations: This criterion covers the assessment of a vendor’s employee retention rates and commitment to corporate social responsibility (CSR) initiatives, including data for good.

Ability to Execute Evaluation Criteria

<i>Evaluation Criteria</i>	<i>Weighting</i>
Product or Service	High
Overall Viability	High
Sales Execution/Pricing	Low
Market Responsiveness/Record	Medium
Marketing Execution	Medium
Customer Experience	Medium
Operations	Low

Source: Gartner (June 2026)

Gartner evaluates a vendor's Completeness of Vision in the AI platforms for DSML market by applying criteria that assess its ability to understand current market trends across diverse user groups, including data scientists, AI engineers, software developers, and business users. This involves recognizing how emerging technologies can advance the goals of data-driven decision making and the development and deployment of AI systems.

Completeness of Vision

Market Understanding: This criterion covers the assessment of a vendor's vision and strategic positioning with respect to GenAI and the DSML market as a whole, the competitive differentiators of the vendor, and the awareness of customer needs, buying personas, and end-user roles.

Marketing Strategy: This criterion covers the assessment of a vendor's go-to-market approach, considering targeted buyer and user personas and market visibility in the context of dedicated resources.

Sales Strategy: This criterion covers the assessment of a vendor's ability to produce different sales strategies for target customer personas, strategic partnerships, and effective use of sales resources.

Offering (Product) Strategy: This criterion covers the assessment of a vendor's depth and breadth of product portfolio and its open-source strategy and commitment.

Business Model: This criterion covers the assessment of a vendor's evolution of its business model in line with enterprise needs, the significance of DSML to its overall business, and the viability of the vendor in the context of current market dynamics.

Vertical/Industry Strategy: This criterion covers the assessment of a vendor's clear and consistent vertical focus and understanding of industry needs.

Innovation: This criterion covers the assessment of a vendor's GenAI feature quality and roadmap, and commitment to ongoing research that results in significant product enhancements.

Geographic Strategy: This criterion covers the assessment of a vendor’s clear and consistent regional focus and understanding of regional market dynamics.

Completeness of Vision Evaluation Criteria

<i>Evaluation Criteria</i>	<i>Weighting</i>
Market Understanding	High
Marketing Strategy	Medium
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	Low
Innovation	High
Geographic Strategy	Low

Source: Gartner (June 2026)

Quadrant Descriptions

Leaders

Leaders in this market have a mature, refined, and targeted company and platform strategy that incorporates and leverages GenAI and AI agents to drive their customers’ business value. They see opportunities for leveraging agents that other providers may not see or have made significant investments above and beyond standard offerings. They have the capability to innovate at a speed that outperforms other vendors. In addition,

they can clearly articulate how they provide value to the multiple types of personas involved in the process of building data science and machine learning models.

Challengers

Challengers in this market have the operational capacity to serve a wide variety of enterprise needs in the data science and machine learning space through brand recognition and complementary product offerings. Their current limitations are centered on the appeal of the platform among their target users and limited geographic and/or industry focus. They show potential for adding innovative and differentiating features from their product roadmap, which could gain traction in specific industries or use cases that are in high demand.

Visionaries

Visionaries understand the DSML market and its future direction, and offer a differentiated view of solutions that need to be provided to meet enterprise AI needs. They offer industry-specific functionality and clearly demonstrate value to their customers on an individual and enterprise level. They are limited by not having the necessary recognition of their product for complete end-to-end DSML capabilities due to historical brand association or limited marketing initiatives and community influence.

Niche Players

Niche Players are characterized by their specialized focus on specific industries or a limited set of user groups, such as low-code or code-first practitioners. They are known to deliver solutions that meet the needs of their target demographics but fail to demonstrate a broader appreciation of market trends and enterprise needs, especially in regard to GenAI. Their appeal is limited beyond a core audience, and they struggle to grow in line with the average market rate. They struggle with user adoption of the platform beyond a basic set of use cases, such as data preparation and exploration, that do not drive as much business value as competitor offerings.

Context

The AI platform for data science and machine learning (DSML) market remains central for enterprise AI delivery, evolving beyond classical data science techniques into holistic

platforms supporting end-to-end AI model and agent development. These platforms now provide comprehensive capabilities for modern generative AI solutions, including retrieval-augmented generation (RAG), large reasoning models (LRMs), and sophisticated AI agents built by data scientists and AI engineers. By combining traditional machine learning and optimization with GenAI, enterprises can process multimodal data to automate diverse business processes like predictive maintenance, recommendations, and advanced search. These platforms also enable the creation of AI agents that drive autonomous operations within AI-powered workflows and perform specialized tasks across the DSML life cycle.

As DSML activities expand beyond centralized core teams, platforms increasingly cater to diverse, persona-based user groups, including data analysts, domain experts, citizen data scientists, and business users. This democratization is facilitated through low-code interfaces, natural language tools, and platform-provided AI assistance. To manage these decentralized efforts safely, enterprises must rely on platforms that enforce robust operations practices, implement strict guardrails, and provide AI governance for risk management. AI and analytics leaders must evaluate how the AI platform for DSML supports consistent, reproducible environments, ensuring proper data lineage, explainability, and the continuous monitoring of both models and autonomous agents to deliver secure and scalable business value.

Market Overview

The AI platforms for data science and machine learning (DSML) market is increasingly focused on supporting end-to-end AI models and agent development using a diverse array of data science and AI techniques. GenAI remains a significant catalyst, with modern platforms expanding to incorporate large language models (LLMs) and large reasoning models (LRMs) through provisioning, fine-tuning, and prompt monitoring. The strategic value of DSML platforms within enterprises continues to rise as they tackle the complex challenge of integrating data, code, and infrastructure to build robust, multimodal applications like computer vision and natural language processing. Furthermore, composite AI systems that combine traditional machine learning with techniques like optimization, alongside retrieval-augmented generation (RAG) capabilities, have become standard for developing scalable, advanced solutions.

DSML platforms must efficiently handle various structured, semistructured, and unstructured data across both fully managed cloud services and on-premises infrastructure. As data engineering, model training, and deployment pipelines mature, platforms increasingly cater to specialized AI experts — including data scientists and AI engineers — by providing consistent, reproducible development environments that improve productivity.

The Gartner 2025 Head of Analytics and Data Science Survey showed that 59% of respondents believed their automation of DSML models to be “effective” versus 41% who stated this was “neutral” or “ineffective.”¹ This is likely due to delivery practices not keeping up with technology advancements as multiple Ops (e.g., MLOps, DataOps, ModelOps, AgentOps, and others) practices need to be integrated for deploying models and agents in production. The need for platforms to develop AI agents to meet a diverse set of use cases is an increasing priority for D&A leaders. Choosing between platforms is difficult due to the nature of agents that can be built, ranging from predictable workflows to open-ended goal-seeking. In addition, the target user persona of platforms can vary from citizen developers to AI experts. AI platforms for DSML have been able to provide agent building capabilities without LLMs for some time, which makes them suitable for more advanced use cases where composite techniques and custom tuning of models for agent tasks need to be performed.

The movement to make data science accessible to a broader range of business roles also continues to accelerate. Platforms increasingly empower domain experts, data analysts, and business users through low-code interfaces, natural language tools, and automated AI assistance. A significant challenge for data science leaders is overseeing these decentralized efforts; therefore, enhanced AI governance, guardrails for risk management, and comprehensive data lineage and metadata tracking are now mandatory capabilities to ensure safe collaboration and asset reuse. Vendors have ample opportunity to succeed by providing project management tools that foster collaboration across dispersed teams, which is critical for advancing enterprisewide GenAI initiatives. To successfully deploy these sophisticated systems, platforms must integrate multiple Ops practices capable of orchestrating both batch and real-time workloads.

Developing AI agents to meet diverse use cases is a top priority, with platforms now supporting the creation of LLM-based agents by actively managing models, tools, memory, and protocols. These AI agents can drive autonomous operations within business workflows and perform automated tasks across the entire DSML life cycle. By

unifying classic predictive techniques with modern generative AI, these platforms enable customized, transparent, and continuously monitored models and agents suitable for a wide spectrum of user personas.

Acronym Key and Glossary Terms

AgentOps	AI agent operations
ASI	artificial superintelligence
AutoML	automated machine learning
DSML	data science and machine learning
DataOps	data operations
FDE	forward-deployed engineer
GenAI	generative AI
LLM	large language model
LRM	large reasoning model
MCP	Model Context Protocol
ML	machine learning
MLOps	machine learning operations
ModelOps	model operations

RAG	retrieval-augmented generation
-----	--------------------------------

⊕ Evidence

⊕ Evaluation Criteria Definitions

Unlock more AI value

With unmatched, proprietary AI insights from Gartner

[Explore the AI Hub ↗](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

