

Magic Quadrant for Data Science and Machine Learning Platforms

17 June 2024 - ID G00799982 - 39 min read

By Afraz Jaffri, Aura Popa, [and 5 more](#)

Data science and machine learning platforms allow insight-driven decision making based on data science techniques via build, customize and deploy machine learning and generative AI models. This research helps data and analytics leaders make critical decisions selecting a DSML platform vendor.

Strategic Planning Assumption

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

Market Definition/Description

Gartner defines a data science and machine learning platform as an integrated set of code-based libraries and low-code tooling that support the independent use by, and collaboration between, data scientists and their business and IT counterparts through all stages of the data science life cycle. These stages include business understanding, data access and preparation, experimentation and model creation, and sharing of insights. They also support machine learning engineering workflows including creation of data, feature, deployment and testing pipelines. The platforms are provided via desktop client or browser with supporting compute instances and/or as a fully managed cloud offering.

Data science and machine learning (DSML) platforms are designed to allow a broad range of users to develop and apply a comprehensive set of predictive and prescriptive analytical techniques. Leveraging data from distributed sources, cutting-edge user experience, and native machine learning and generative AI (GenAI) capabilities, these platforms help to augment and automate decision making across an enterprise. They provide a range of proprietary and open-source tools to enable data scientists and domain experts to find patterns in data that can be used to forecast financial metrics, understand customer behavior, predict supply and demand, and many other use cases. Models can be built on all types of data, including tabular, images, video and text for applications that require computer vision or natural language processing.

The supported machine learning techniques range from classic regression or decision trees to more complex deep learning and reinforcement learning and GenAI. The models built using these techniques can be used for tasks within business processes such as credit scoring, churn prediction, predictive maintenance, recommendation and image classification.

Organizations that use DSML platforms can significantly reduce the cycle time and barriers to entry for creating predictive and prescriptive models, generating insights and distributing results. They enable collaboration and reuse of assets across multiple teams and departments and orchestration of workloads to handle high volumes of data. They also provide consistent and reproducible training and development environments that provide lineage between data, code and model assets, enhancing data scientists' productivity. Additionally, low-code and natural language interfaces enable domain experts and business users to create predictive models with simple workflows. DSML platforms support MLOps practices such as deploying models in production for both batch and real-time workloads and provide ongoing monitoring of model metrics and compliance.

Must-Have Capabilities

The must-have capabilities for this market include:

- Import data from databases, data warehouses and file stores located on-premises and in the cloud
- Build and evaluate models using a library of core data science and machine learning techniques, methods, algorithms and processes
- Deploy, host and serve models in the platform for usage in services and applications

Standard Capabilities

The standard capabilities for this market include:

- Ability to build models from structured and unstructured data sources including text, images, video, audio and geospatial
- Low-code interface for model development suitable for nonexpert data science roles, including business users and domain experts
- Notebook-based code interface for data scientists to perform data access, preparation, model development and publication tasks
- Postdeployment model life cycle management to retrain, retire or adapt models based on detecting and analyzing data, feature and model drift
- Support for MLOps-based processes and tools that enable ML models to be deployed at scale in different operational environments

Optional Capabilities

The optional capabilities for this market include:

- Platform-generated recommendations for the best way to prepare, integrate and model data as well as automated creation of machine learning models based on manually selected target prediction
- Advanced interfaces that facilitate more complex modeling for simulation, optimization and deep learning-based use cases
- Custom SDKs that provide more control and flexibility for code-based model development and integration with services and applications
- Functionality for working with GenAI models, such as large language models, through tracking, selection and monitoring of prompts, models and outputs
- Techniques and tools that increase the transparency and interpretability of models to understand how and why model outputs are generated

Magic Quadrant

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



Source: Gartner (June 2024)

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud is a Challenger in this Magic Quadrant. The Alibaba Cloud Platform for AI (PAI) provides high-performance computing resources and infrastructure that data science teams can use to build, deploy, utilize, and manage machine learning (ML) and generative AI (GenAI) models. Its operations are mainly focused in China, Southeast Asia, the Middle East and Australia, and its clients tend to be enterprises of various sizes operating across all

sectors. Alibaba Cloud has made significant investments in graphics processing unit (GPU) availability for GenAI training and access to first- and third-party models for customization and domain-specific use cases.

Strengths

- **Focus on data governance.** DataWorks, Alibaba Cloud's platform as a service (PaaS), provides a leading platform for governance and managing data quality throughout the ML life cycle.
- **Cost-control capabilities.** High-performance computing innovation and cost-management tools enable users to lower AI training costs by optimizing resource utilization.
- **Foundation model availability.** Alibaba Cloud's own set of large language models (LLMs), community and open-source models from partners provide its clients access to various multimodal foundation models.

Cautions

- **Complexity.** With Alibaba Cloud still evolving its PAI's low-code capabilities for model development, high technical competence is required to utilize all platform features.
- **Limited presence outside of China.** Enterprises outside of China may encounter limited support, communities and partners to help with their implementation.
- **Corporate restructuring.** Alibaba Group is undergoing organizational and management changes that raise uncertainty about the future direction of services on Alibaba Cloud.

Altair

Altair is a Leader in this Magic Quadrant. Its RapidMiner platform includes AI Hub, AI Studio, and other Altair products focused on increasing and accelerating access to data insights and predictive models by enabling collaboration between data science and business teams. Its operations are geographically diversified, and its clients tend to be midsize to large enterprises in industries including automotive, aerospace, defense and healthcare. Altair plans to increase democratization of data through investment in providing a conversational interface for workflow authoring using its proprietary analytics translation language. It also addresses enterprise-grade AI by strengthening AI, Internet of Things (IoT) and high-performance computing (HPC) convergence through integration with other Altair products.

Strengths

- **Market understanding.** Altair's RapidMiner platform addresses enterprise pain points for enterprise GenAI and ML adoption, measuring success by completed project deployments rather than platform usage.
- **Industry focus.** RapidMiner's 2022 acquisition by Altair has opened new opportunities, particularly in asset-centric industries and scientific applications.
- **Pricing.** Altair's pricing structure is flexible and not based on named user licenses or consumption levels, but on units tied to usage of the Altair software portfolio.

Cautions

- **Market traction.** Altair is placed in the lower quartile of Gartner's Customer Interest Indicator, indicating a lack of awareness of capabilities among end users.
- **Product coherence.** Integration between rebranded and legacy products with overlapping capabilities in the Altair RapidMiner platform is continuing and will need to be completed over time.
- **Acquisition friction.** The embedding of RapidMiner within the Altair ecosystem can increase end-user demand for support, which could potentially result in lower levels of customer interaction.

Alteryx

Alteryx is a Niche Player in this Magic Quadrant. Its Alteryx Designer, Alteryx Server and Alteryx Analytics Cloud products provide an AI platform for enterprise analytics, and enable analysts, data scientists and business users to use data for decision making. Its operations are geographically diversified, and its clients tend to be enterprises of a range of sizes operating across all sectors. Alteryx is developing a suite of AI tools, under the AiDIN brand, that targets the needs of all personas. In March 2024, the company was acquired by Clearlake Capital Group and Insight Partners.

Strengths

- **Insights-driven vision.** Alteryx has targeted its platform to all roles involved in data science, aiming to increase efficiency through augmentation.
- **User sentiment.** Alteryx has a large, passionate user community and a long history of customer satisfaction.

- **Automation and augmentation.** Alteryx simplifies complex analytics processes and data preparation through AI-generated augmentation of key tasks and automation of end-to-end processes.

Cautions

- **Recent acquisition.** In March 2024, Alteryx was acquired by two private equity companies. The implications of this acquisition on the company's existing partnerships and future direction are still being determined.
- **Competitive differentiation.** Alteryx struggles to stand out in a crowded market where much of the functionality that brought it initial success has been replicated or enhanced.
- **Expert data scientist focus.** Alteryx is known for its low-code platform, to which it has recently added code-centric functionality. Data scientists will need to be convinced the platform can meet their needs.

Amazon Web Services

Amazon Web Services (AWS) is a Leader in this Magic Quadrant. Its SageMaker, Bedrock, SageMaker Canvas and SageMaker Data Wrangler products are broadly focused on enabling enterprises to develop, train and host DSML and GenAI models through code-based and low-code tooling. Its operations are geographically diversified, and its clients consist of enterprises of all sizes and sectors, with specific initiatives for the public sector and startups. AWS is focused on making GenAI enterprise-ready by delivering purpose-built compute infrastructure (including specialized Trainium and Inferentia chips), increasing access to commercial and open-source models, and improving security, governance and privacy controls.

Strengths

- **Security compliance.** AWS provides access to a variety of GenAI foundation model providers with enhanced privacy and security guarantees, including fine-tuning and retrieval-augmented generation (RAG) use cases.
- **Custom AI infrastructure.** AWS has made significant investment in providing custom hardware and services, such as SageMaker HyperPod for training and serving large models at scale.

- **Training and community.** AWS offers a number of in-person and virtual events targeted at different user groups, including re:Invent, Summits and Dev Days.

Cautions

- **Core data science.** AWS has given more attention to ML engineering than decision augmentation-based data science, causing end users to explore additional services to provide core functionality.
- **Complexity.** AWS' large portfolio for data science, ML and GenAI can cause difficulties in identifying the right product for a business problem.
- **GenAI market sentiment.** AWS' own Titan foundation models have limited brand recognition, weakening its position as an innovative model provider.

Anaconda

Anaconda is a Niche Player in this Magic Quadrant. Its Anaconda data science platform enables enterprises to utilize open-source Python libraries and tools in a single platform. Its operations are geographically diversified, and its clients tend to be small, midsize and large enterprises across all sectors. Anaconda supports its own offering with integration with other DSML platforms, including IBM, Databricks and Domino Data Lab. It recently introduced its GenAI-powered Anaconda assistant to both cloud and local deployments of the platform.

Anaconda 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

- **Large user community.** The open-source Anaconda Python Distribution and conda package and environment management tool are well-known and used by thousands of data scientists.
- **Tooling.** Anaconda offers open-source tools that expert data scientists need to provide core data science and machine learning engineering capabilities.
- **Low barrier to entry.** Anaconda offers a free edition of its platform that is well-used, well-maintained and supported by many different groups, significantly lowering the barrier to entry for new practitioners.

Cautions

- **Scope of vision.** Anaconda shows a lack of vision in providing out-of-the-box capabilities for data science tasks, in favor of relying on open-source components.
- **Target personas.** Anaconda targets expert data scientists without addressing the needs of data engineers, MLOps engineers and team leads.
- **Differentiation.** The commercial offering of the platform does not offer significant benefits for users, as package management and DSML libraries are available from other vendors in the market.

Cloudera

Cloudera is a Visionary in this Magic Quadrant. Its Cloudera Data Platform and Cloudera Machine Learning products provide a unified, scalable, and open data, analytics and AI platform for data engineering and machine learning workloads, with end-to-end governance and model hosting capabilities. Its operations are geographically diversified, and its customers tend to be large enterprises across all sectors. Cloudera is focused on embedding GenAI capabilities into its platform through AI assistants for code and data visualization generation, and enabling users to build and deploy GenAI applications from a variety of model providers.

Strengths

- **Hybrid and multicloud.** Cloudera provides an integrated data and DSML platform built for hybrid cloud, multicloud and on-premises environments, and complex workloads.
- **Secure open source.** Cloudera leverages popular open-source components that are familiar to data scientists, together with enterprise-level security that minimizes lock-in.
- **Use-case templates.** Cloudera offers a wide selection of prebuilt GenAI and DSML use-case templates for a variety of foundation models and techniques.

Cautions

- **Complexity.** Although low-code functionality, such as augmented assistance and visual workflows for business and domain-focused personas, is improving, it is limited compared with other vendors in this Magic Quadrant.
- **Partnerships.** Compared with vendors in this Magic Quadrant, Cloudera has fewer stable partnerships with systems integrators that have dedicated resources to help with

enterprise deployments.

- **Market position.** Cloudera is still developing the vision for its platform to move the company from being known as a data vendor and toward being also known as an AI vendor.

Databricks

Databricks is a Leader in this Magic Quadrant. Its Databricks Data Intelligence Platform combines data management, data warehousing, governance and engineering with data science, ML and AI development for data and AI workloads. Its operations are geographically diversified, and its clients tend to be midsize and large enterprises across all sectors. The company recently acquired MosaicML (July 2023) and Einblick (January 2024) to drive GenAI features in its platform and improve support for low-code development, respectively.

Strengths

- **Unified platform.** Databricks' platform combines a lakehouse platform and centralized governance with DSML functionality to simplify data and ML engineering team processes and management.
- **Acquisitions.** Recent acquisitions, including MosaicML, enable Databricks' users to utilize large language models quickly and at lower cost than competitors.
- **Future-facing strategy.** Databricks has invested in its vision of the future of user demands by bringing GenAI and low-code capabilities to its platform.

Cautions

- **Performance.** Gartner Peer Insights reviews show some users are unhappy with the performance tuning and cluster management required when developing and deploying DSML workloads.
- **Upgrade friction.** Databricks' rapid product expansion leads to new product versions that can require customers to keep up with updates and ensure stability of existing workloads.
- **Solution alignment.** Databricks' all-in-one solution and lakehouse architecture requires additional design considerations when used with composable, hybrid data architectures.

Dataiku

Dataiku is a Leader in this Magic Quadrant. Its Dataiku product is focused on providing a collaborative experience of low-code and code-based development for data scientists and business users, with augmented insights to accelerate model building and deployment. Its operations are geographically diversified, and its customers tend to be midsize and large enterprises across all sectors. Dataiku leads the LLM Mesh initiative to democratize GenAI development and is expected to expand the ecosystem with other partners.

Strengths

- **Collaboration.** Dataiku's platform is designed to enable diverse roles to work together to operationalize AI through bringing together code and low-code development.
- **GenAI vision.** Dataiku, along with partners, founded the LLM Mesh initiative focused on providing enterprisewide usage of GenAI models to abstract away the complexity of managing cost and performance.
- **People focus.** Dataiku's focus on the impact of AI on people and business through change management gives AI leaders confidence to utilize the technology enterprisewide.

Cautions

- **User community.** Dataiku's user community is not as vast, accessible or at the same level of interaction as those of other vendors in this Magic Quadrant, limiting the number of support channels available to users.
- **Pricing.** Dataiku has a complex pricing model based on platform fees and per-user licenses, formed around four types of personas.
- **Administration.** Maintaining the on-premises offering of the platform can be resource-intensive, especially for upgrading and resource management.

DataRobot

DataRobot is a Leader in this Magic Quadrant. Its DataRobot AI platform focuses on accelerating the time to value for data science and ML projects through advanced automated ML capabilities, code-based development and governance, and monitoring for ML models. Its operations are geographically diversified, with the exception of China, and its clients comprise enterprises of a variety of sizes and sectors. The company has increased its focus on governance and assurance for ML and GenAI models through enhanced platform

capabilities and participation in the U.S. Department of Commerce's U.S. AI Safety Institute Consortium.

Strengths

- **Ease of use.** DataRobot's abstraction of GenAI and predictive model building simplifies processes for data scientists and business users.
- **Focus on value generation.** DataRobot uses targeted customer-facing teams for customer engineering and applied AI to accelerate value delivery.
- **Market understanding.** DataRobot's strategy focuses on collaboration between multiple participants in end-to-end DSML activities, enabling business-focused personas to ensure their needs are addressed.

Cautions

- **Operational practices.** DataRobot has experienced changes in its leadership team and high staff turnover, which may hamper the company's efforts to maintain customer satisfaction.
- **Pricing.** DataRobot's pricing model and license cost are among the highest in this report.
- **Product strategy.** DataRobot promotes itself as an AI platform but its focus on data management, visualization, exploration, annotation and labeling are behind those of peers in the market.

Domino Data Lab

Domino Data Lab is a Visionary in this Magic Quadrant. Its Domino enterprise AI platform focuses on flexibility by allowing data scientists to use on-premises and cloud infrastructure, as well as a variety of development environments and languages, for development, deployment and governance of DSML models. Its operations are geographically diversified, and its clients tend to be global enterprises in the life sciences, financial services, public sector and regulated industries. The vendor continues to broaden its appeal to include a wider segment of end users by adding GenAI functionality and improving trustworthiness with responsible AI capabilities.

Strengths

- **Security.** Domino Data Lab differentiates itself in the market by focusing on verticals that require increased levels of security.

- **Governance.** Domino Data Lab's governance of data science and machine learning assets enables the vendor's focus on regulated industry verticals.
- **Flexibility.** Extensive compatibility with developer tools, languages and development environments makes Domino Data Lab's platform an option for enterprisewide DSML deployment.

Cautions

- **Market penetration.** Two-thirds of Domino Data Lab's customers are concentrated in four industries, limiting its potential to appeal to a broad range of enterprises.
- **Complexity.** Domino Data Lab is focused on serving large, expert data science teams and is the best option for enterprises with purely line-of-business (LOB) user needs.
- **Low code.** Domino Data Lab offers low-code options, but it lags behind the offerings of other vendors in this Magic Quadrant.

Google

Google is a Leader in this Magic Quadrant. Its Vertex AI platform is focused on enabling data scientists, machine learning engineers, and business users to build and deploy DSML models using Google Cloud Platform (GCP) infrastructure. Its operations are geographically diversified, with the exception of China, and its customers are enterprises of all sizes across a wide variety of industries. Google is invested in advancing GenAI research and development with its Gemini family of foundation models produced by Google DeepMind.

Strengths

- **Foundation models.** Foundation models available from Google's in-house R&D and third-party providers, for both commercial and open use, ensure that end users benefit from cutting-edge innovation.
- **Balanced portfolio.** Investment in GenAI has not derailed Google's attention from enhancing existing DSML methods and capabilities in its Vertex AI platform.
- **Delivery.** Google rolled out more than 500 updates to Vertex AI in 2023, further enhancing its positioning as a platform centered on innovation and meeting end-user needs.

Cautions

- **Data and AI governance.** Unification of data and AI governance capabilities lags behind competitors.
- **Platform ecosystem.** Vertex AI makes sense as a DSML platform for enterprises that have invested or will invest in Google Cloud Platform as their data storage and processing platform of choice.
- **Core data science.** Google's focus on core data science use cases is weak compared with its focus on machine learning.

H2O.ai

H2O.ai is a Visionary in this Magic Quadrant. Its H2O Driverless AI, H2O Hydrogen Torch, H2O Document AI, H2O AI Feature Store and H2O MLOps products are focused on democratizing AI through open-source offerings and AI-augmented tooling. Its operations are geographically diversified, and its clients tend to be enterprises of varying sizes, mainly in the financial services, telecommunications and consumer packaged goods industries. H2O.ai plans to invest in expanding its product lines and feature sets to make use of GenAI models and techniques, and to co-create innovations with clients using its own open-source foundation model.

Strengths

- **Customer collaboration.** H2O.ai works closely with customers to identify, remediate and co-create product functionality.
- **Generative AI.** H2O.ai has trained its own open-source foundation model and created a GenAI ecosystem for different foundation models, vector databases, hardware and hosting.
- **Open-source community.** A large community of data scientists uses open-source H2O.ai, strengthening the vendor's understanding of data scientists' needs as well as key trends and pain points.

Cautions

- **Dependence on partners.** H2O.ai leverages partnerships for some DSML features, such as data management, data preparation and governance.
- **Limited presence outside of key markets.** H2O.ai has a presence in telecommunications and consumer packaged goods industries, but most customers are concentrated in

banking, financial services and the public sector.

- **Product portfolio coherence.** The division of functionality between products means the ability to collaborate between multiple personas and stakeholders is limited.

IBM

IBM is a Challenger in this Magic Quadrant. Its watsonx product suite provides an open platform that delivers business AI use cases integrating data management, governance and GenAI end-to-end tooling for data scientists and developers. Its operations are geographically diversified, and its customers comprise enterprises of a variety of sizes and sectors. The company is offering its Granite series of foundation models with generalized and domain-specific versions that provide cost-efficient alternatives to larger commercial models.

Strengths

- **Investment in responsible AI.** IBM has partnered with Meta to co-launch the AI Alliance to support open science and open innovation in AI.
- **Openness and flexibility.** IBM's tight integration with Red Hat OpenShift, and partnership ecosystem, provides adaptability and flexibility for enterprise needs in multiple deployment scenarios.
- **Market understanding.** IBM has targeted governance of AI as a market need, aligning with the increased demand from enterprises with security and compliance requirements.

Cautions

- **Market traction.** The adoption of watsonx is typically seen among customers who already have a significant investment in other IBM products.
- **Product coherence.** watsonx is the latest marketing name for the Watson brand. watsonx itself consists of three core components, which could cause confusion for buyers.
- **Pricing structure.** Gartner clients report difficulty navigating IBM and watsonx's complex pricing.

KNIME

KNIME is a Niche Player in this Magic Quadrant. Its KNIME Analytics Platform and KNIME Business Hub serve business users, domain experts, analysts, and data scientists with a

flexible and extendable open-source platform for end-to-end DSML tasks. Its operations are focused in Europe and the U.S., with presence in other regions through its partner network. Its clients comprise enterprises of a variety of sizes and sectors. The vendor has delivered GenAI-based assistants, known as K-AI, for data exploration and low-code and code-based activities. It also recently offered the Community Hub version of its software, in free and paid editions, as a fully managed SaaS platform.

Strengths

- **Democratization.** KNIME's desktop analytics platform is available as a free and unrestricted download, and it enables users that have a range of analytical skill sets, as well as those users getting started in the data science field.
- **Data science depth.** KNIME's range of out-of-the-box nodes and components covers the entire spectrum of DSML techniques, from simulation to econometrics and deep learning.
- **Use cases.** KNIME's community of more than 400,000 users has contributed a wide variety of industry and horizontal solutions, in addition to those from KNIME and its partners.

Cautions

- **Global reach.** KNIME has direct presence in Europe and the U.S., and uses partners to serve clients in other regions.
- **Market traction.** KNIME is placed in the lower quartile of Gartner's Customer Interest Indicator, indicating a lack of awareness of capabilities among end users.
- **Target personas.** KNIME's strategy focuses on enabling all types of users through low-code development, but code-first development needs are underserved.

MathWorks

MathWorks is a Niche Player in this Magic Quadrant. Its MATLAB and Simulink products provide engineering teams with a scalable platform for DSML model building and deployment, as well as a suite of simulation tools for industrial applications. Its operations are geographically diversified, and its clients tend to be enterprises of all sizes in industrial sectors such as manufacturing, telecommunications and life sciences. MathWorks' current focus is on an AI Chat Playground for testing and evaluating the output MATLAB code

responses from natural language instructions for different LLMs. MathWorks provides specific toolboxes for several engineering personas.

Strengths

- **Longevity.** MathWorks is a longstanding vendor that provides a stable set of core DSML tools and continues to evolve with market trends.
- **Embedded AI.** MathWorks allows engineers to embed AI into industrial devices, appliances and IoT assets without needing specialist deployment knowledge.
- **Safety-critical systems.** MathWorks is ubiquitous in sectors where safety is paramount, such as the aerospace and automotive industries, providing a suite of tools for robust DSML development.

Cautions

- **Innovation.** MathWorks does not operate at the cutting edge of GenAI, and many use cases outside of asset-centric industries may not be as well-served.
- **Complexity.** Despite the presence of some low-code tooling, MathWorks is not a toolset for newcomers to DSML development.
- **Trajectory of vision.** MathWorks continues to target technical LOB users, such as engineers, limiting its usage in other business domains.

Microsoft

Microsoft is a Leader in this Magic Quadrant. Its Azure Machine Learning (Azure ML) products provide data scientists and LOB, low-code developers with infrastructure, models and tooling to build predictive and GenAI models that can be scaled across the enterprise. Its operations are geographically diversified, and its customers comprise enterprises of a variety of sizes and sectors. Microsoft is focused on offering access to foundation models from its internal research group and its partnership with OpenAI and other providers. It also has invested in advanced data centers with custom hardware and chips for GenAI usage.

Strengths

- **Enterprise GenAI.** Microsoft has continued to innovate in its offering of foundation models through its Azure ML model catalog by providing a range of models, deployment options and pricing options.

- **Upskilling.** Microsoft produces an extensive range of training resources aimed at beginners, experts and executives in the DSML space, including new Azure-based GenAI certification.
- **Research and development.** Microsoft's own research group has contributed with its own smaller foundation models, prompt engineering frameworks and new architectures for retrieval-augmented generation (RAG) implementation.

Cautions

- **Separate product lines.** Frequent rebranding of components and separate product lines within Azure ML create friction for data science teams looking for a unified platform that supports collaboration.
- **Data exploration and visualization.** Exploratory data visualization capabilities in Azure ML are provided through Power BI. Power BI's offering through Microsoft Fabric will trigger reevaluation for many enterprises.
- **Copilot integration.** GitHub Copilot requires additional assessment for full functionality across Azure ML compute instances.

Posit

Posit is a Niche Player in this Magic Quadrant. Its Posit Workbench, Posit Connect, Posit Package Manager and Posit Cloud products provide an open platform for data scientists to build predictive and prescriptive models to deliver insights to business users and decision makers. Its operations are focused in the U.S., with operations in other regions offered via partners, and its customers tend to be small, midsize and large enterprises across all industries. Posit is focused on adding features to support data scientists, such as streamlining management of code and applications. The company recently announced a partnership with Databricks for enhanced integration.

Strengths

- **Corporate citizenship.** Posit's corporate social responsibility commitments include community initiatives and programs for public good written into its charter, aligning with many corporate ESG objectives.
- **Partnerships.** Expanding partnerships with major cloud and data platforms will make integration and usage of enterprise data with its platform easier.

- **Dual usage.** Posit's products can function as a low-cost DSML, analytics and business intelligence (ABI) platform instead of separate solutions for data science teams and analysts.

Cautions

- **ML engineering.** Posit's focus is on insight delivery and not on elements of ML engineering, such as deployment, monitoring and governance functionality.
- **Low code.** With limited low-code functionality, Posit's platform is geared toward data scientists, with use for business analysts restricted to dashboard or application usage.
- **GenAI.** Posit lacks a clear roadmap for GenAI in areas such as foundation model development, fine-tuning LLMs and agent development.

SAS

SAS is a Leader in this Magic Quadrant. Its SAS Viya product focuses on providing a scalable DSML platform suited to a wide range of industry use cases, with proprietary algorithms, coding language and models, as well as open-source integration. Its operations are geographically diversified, and its customers tend to be midsize to large enterprises across a variety of industries. SAS plans to add to its industry and value focus by leveraging its relationship with Microsoft to bring more GenAI capabilities into the product.

Strengths

- **Partner network.** SAS's network of more than 1,700 partners across technology, resellers and services leads to greater adoption of its platform, with a global reach.
- **Strong industry focus.** SAS's customized models, workflows and solutions reduce the time to value for a variety of customer departments and business units.
- **DecisionOps.** SAS is one of the few providers to focus on operationalization of decisions as well as models, accelerating business value through process optimization.

Cautions

- **Cost.** Discussions with users of Gartner's client inquiry service, together with evidence from Gartner Peer Insights, indicate the SAS platform is perceived as expensive compared with competitors' offerings. As the company works toward an initial public offering (IPO), further pricing amendments can be expected.

- **SAS language.** SAS has added open-source support, but the decline in usage of SAS as a programming language reduces the attractiveness of the product compared with other platforms.
- **GenAI functionality.** SAS has not been an early mover with GenAI, and many GenAI functions embedded in the platform are in preview form.

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

As this is a new Magic Quadrant, no vendors were added.

Dropped

As this is a new Magic Quadrant, no vendors were dropped.

Inclusion and Exclusion Criteria

To qualify for inclusion, providers need to:

- Demonstrate a go-to-market strategy for their DSML platform that focuses on the persona of a professional data scientist.
- Support the following use case:
 - Expert data science teams — Use of the DSML platform by a team of expert data scientists in a central or distributed function utilizing a primarily code-driven environment.

In addition, the DSML platform must support at least one of these use cases:

- Line-of-business teams — Use of the DSML platform by a team of business users (i.e., citizen data scientists) in one or more business function(s), primarily through low-code development.
- Fusion teams — Use of the DSML platform by a multidisciplinary team that blends information technology, data science, and business domain expertise using code and low-code development.

International presence:

- The vendor must have at least 10 paying customers (logos) as of 1 December 2023 for their DSML platform in each of three of the following regions:
- North America
- South America
- Europe, the Middle East and Africa
- China
- Central Asia
- Asia/Pacific

Vertical markets:

- The vendor must have at least 10 paying customers (logos) as of 1 December 2023 for their DSML platform in each of four of the following vertical/industry categories:
- Banking and financial services
- Communications, media and entertainment
- Education
- Government
- Healthcare provision and life sciences
- Manufacturing and natural resources
- Retail

- Logistics and transportation
- Utilities
- Technology

Customer Interest Index:

- The vendor must rank among the top 20 DSML platforms for this Magic Quadrant according to Gartner's Customer Interest Index (CII). Data inputs used to calculate relevancy include interest from Gartner clients, customer reviews and key vendor growth metrics.

Honorable Mentions

Baidu

Baidu AI Cloud is a popular cloud vendor in China. The Baidu Machine Learning (BML) product provides an end-to-end platform for DSML for both code-centric and low-code development. Key features include a custom data processing engine, together with object storage, model and feature management, and efficient hardware for training and serving ML models, including its PaddlePaddle deep learning library for advanced model building. Baidu did not meet the geographic coverage criteria to be included in this Magic Quadrant.

Oracle

Oracle has added DSML capabilities to its Oracle Cloud Infrastructure (OCI) platform by offering OCI Data Science and OCI Generative AI service. OCI Data Science is a fully managed platform built for expert data scientists to utilize open-source tools and specialized hardware for model building and operationalization. OCI Generative AI is also a fully managed service for GenAI development, including access to leading foundation models and fine-tuning capabilities. Oracle did not meet the use-case criteria for LOB and collaborative team usage for this Magic Quadrant.

Salesforce

Salesforce has a suite of AI tools focused on increasing user productivity within the platform but has also added capability to leverage data from outside the platform and utilize GenAI models for custom use cases. Einstein 1 Studio incorporates a set of low-code tools to enable admins and developers to customize AI for their CRM, including Einstein Copilot.

Platform features include Copilot Builder for extending Einstein Copilot with custom actions and Model Builder for building new ML models on Data Cloud. Additional features also include importing predictive and generative models, such as LLMs from OpenAI and Azure OpenAI that a customer may already own, as well as Prompt Builder for prompt development and activation. Salesforce did not meet the criteria for offering a go-to-market strategy for the professional data scientist.

Snowflake

Snowflake is a cloud-native data platform with an architecture built to run analytical, AI and app workloads on a common data foundation, working against multiple data types and programming languages. Using the Snowpark ML Python library, teams can interact with platform components for end-to-end ML model development and operations. Snowflake also natively hosts Streamlit, a popular library for developing interactive AI applications. Snowflake did not meet the inclusion criteria for being in the top 20 providers in the Customer Interest Index for this Magic Quadrant.

Evaluation Criteria

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 machine learning 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 a vendor's customers' willingness to make strategic investments and the 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 diversity, equity and inclusion (DEI).

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	Low
Marketing Execution	Medium
Customer Experience	Medium
Operations	Low

Source: Gartner (June 2024)

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

<i>Evaluation Criteria</i>	<i>Weighting</i>
Business Model	Medium
Vertical/Industry Strategy	Low
Innovation	High

Source: Gartner (June 2024)

Quadrant Descriptions

Leaders

Leaders in this market have a mature, refined, and targeted company and platform strategy that incorporates and leverages GenAI to drive their customers’ business value. They see opportunities for leveraging GenAI 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 complementing product offerings. Their current limitations are centered on the appeal of the platform’s usability to code-first data scientists, ML engineers and analysts using low-code tools. They show potential for adding innovative and differentiating features from their product roadmap that 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 go deep into demonstrating 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

Innovation in the DSML market, like so many others, has been dominated by giving enterprises the tools they need to succeed with GenAI. The trajectory of this market, however, is different due to the dual function of the platforms in relation to GenAI. First, there is the accelerating development of data scientists and low coders through natural language instructions and AI assistants. Second, and more important, there is the need to be able to access, utilize and customize foundational models for enterprise needs. In particular, the MLOps side of GenAI, and all DSML models for security, durability and governance, is given particular emphasis. Use of GenAI in the enterprise will also require other types of models built with various techniques to be added in a composite fashion, which is a key differentiator for many of the platforms.

DSML activities within enterprises have grown outside of centralized, core DSML teams. This results in many enterprises needing to deal with multiple platforms while maintaining best practices for development, monitoring and responsible AI use. Many of the vendors in this Magic Quadrant offer partnerships and integration with each other, meaning that platform selection sometimes is not simply a straight choice between platforms. AI and analytics leaders should understand how these platforms complement each other, whether for low-code, code-first, data management or operationalization-based functionality, to provide the best means of delivering value.

Market Overview

The importance of DSML platforms as a strategic enterprise asset has never been greater. The surge in demand for AI solutions, including GenAI, is at its peak, yet the raw materials of data, models, code and infrastructure have never been more complex to assemble into trusted, scalable products. In the DSML Customer Reference Survey for this Magic Quadrant, 53% of respondents cited GenAI demand as driving a major increase in DSML platform spend in 2024 and beyond.

DSML platforms have evolved to cover full-stack data science, from infrastructure utilization in multicloud environments to data engineering and pipeline building, model training and experiment tracking, deployment and model tracking, and front-end applications. Differentiation between platforms is now based on the amount of abstraction that is provided so team members can quickly iterate on solutions without needing to dive deep into each layer.

GenAI has also accelerated the trend for democratization of data science to support personas more closely aligned with business units and functional areas. AutoML has been a well-established product feature and is now supported with coding assistants, natural language querying and workflow generation. The issue for data science and AI leaders is how to manage and provide governance over the activities of distributed DSML teams and maximize efficiencies through collaboration with centralized resources.

Due to the availability of compute, data and infrastructure needed for DSML development, hyperscalers' offerings are gaining more traction in the market. Yet, there is still room for others to thrive, especially in relation to enabling collaboration between teams, a key pillar for DSML and GenAI development. Bringing DSML techniques to more enterprises, and every area of the enterprise, is an opportunity that can be grasped by vendors and end users alike. The foundational use case of data science for insight-driven decision making must not be lost in the GenAI noise, and DSML platforms offer the perfect place to unite advanced analytics and AI development.

⊕ Evidence

⊕ Evaluation Criteria Definitions

© 2025 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.

[About](#) [Careers](#) [Newsroom](#) [Policies](#) [Site Index](#) [IT Glossary](#) [Gartner Blog Network](#) [Contact](#) [Send Feedback](#)



© 2025 Gartner, Inc. and/or its Affiliates. All Rights Reserved.