



NUCLEUS
RESEARCH

DSML Platform Technology Value Matrix 2026

ANALYST

Duncan Van Kouteren

The Bottom Line

In 2026, the DSML platform market is defined by the emergence of agentic AI as the central organizing principle of enterprise data science and machine learning strategy. Every platform evaluated in this year's Value Matrix ships agent capabilities in some form, moving the market beyond platforms that assist users in building models to platforms that deploy autonomous agents capable of executing multi-step analytical workflows with minimal human oversight. The adoption of Model Context Protocol (MCP) as an interoperability standard reflects a market recognition that the value of agentic AI scales with connectivity, enabling agents to interact with enterprise tools, data sources, and other agents across platform boundaries. Governance capabilities mature in parallel, with platforms providing human-in-the-loop controls, agent monitoring, prompt injection guardrails, and audit trails that capture the full reasoning chain of autonomous actions. For practitioners, agents now automate feature discovery, hyperparameter optimization, model architecture evaluation, and retraining pipelines, shifting the data scientist's role from hands-on model building to orchestrating and refining what agents produce. Agent-driven interfaces simultaneously expand platform access to business analysts and domain experts through natural language. This convergence of agentic execution, open interoperability, and maturing governance redefines the DSML platform as the orchestration layer where data pipelines, model lifecycle management, and autonomous decision-making converge under enterprise-grade control.

Market Overview

The DSML platform market in 2026 is defined by the emergence of agentic AI as the central organizing principle of enterprise technology strategy. Every platform evaluated in this year's Value Matrix has shipped agent capabilities in some form, whether through dedicated agent development frameworks, visual agent builders, or managed agent runtime infrastructure. This shift marks a fundamental evolution from platforms that assist users in building models to platforms that deploy autonomous agents capable of executing complex, multi-step analytical workflows with minimal human oversight. The convergence of large language models, retrieval-augmented generation (RAG), and tool-use capabilities has created a new class of enterprise AI system: one that reasons, plans, retrieves context, and takes action across data pipelines, model development, and operational decision-making.

The defining shift in the DSML market in 2026 is the transition from AI-assisted analytics to autonomous, agent-driven workflows that execute complex multi-step processes with minimal human intervention.



A defining development in 2026 is the emergence of Model Context Protocol (MCP), an open standard for connecting AI agents to enterprise tools and data sources, as an interoperability framework for agentic systems. Multiple vendors have adopted MCP to enable AI agents to discover and interact with enterprise tools, data sources, and other agents across platform boundaries. This standardization reflects a market recognition that the value of agentic AI scales with connectivity. Agents that can access only a single platform's tools are inherently limited compared to those that operate across an organization's full technology stack.

Governance capabilities have matured in parallel, addressing the specific control and compliance challenges introduced by agentic deployments. Platforms now provide configurable human-in-the-loop controls, agent monitoring with real-time performance metrics, guardrails for prompt injection and content moderation, and audit trails that capture the full reasoning chain of autonomous agent actions. As organizations transition from experimental pilots to production-scale agentic deployments, the ability to maintain control, transparency, and accountability over autonomous systems has emerged as a primary differentiator among platforms.

These developments collectively redefine both the evaluation criteria for DSML platforms and the role of the practitioners who use them. The platform is no longer assessed solely on its ability to support model development and deployment; it is now the orchestration layer where data pipelines, model lifecycle management, agent execution, and governance converge. For data scientists, agentic capabilities are reshaping how core platform tasks are executed. Agents now automate feature discovery across large datasets, run hyperparameter optimization sweeps, evaluate candidate model architectures, and trigger retraining pipelines when performance drift is detected. This shifts the data scientist's role from hands-on model building to orchestrating, validating, and refining what agents produce across the model lifecycle. At the same time, agent-driven interfaces are expanding platform access to business analysts and domain experts who can now initiate tasks like demand forecasting, customer segmentation, or pipeline diagnostics through natural language, work that previously required specialized data science skills. Agents inherit the strengths and limitations of the platforms on which they operate, meaning that data quality, feature engineering rigor, and model management discipline are more consequential in an agentic context, not less. The vendors that lead in this year's Value Matrix are those that have built agentic capabilities on strong DSML foundations while

Model Context Protocol has emerged as a critical interoperability standard, enabling AI agents to discover and interact with enterprise tools, data sources, and other agents across platform boundaries.

delivering the connectivity, transparency, and control that production-scale autonomy demands.

Leaders

Leaders in the DSML Platform Technology Value Matrix include Alteryx, Amazon, Databricks, Dataiku, SAS, and Teradata.

Alteryx

Alteryx is recognized as a leader in this year's DSML Platform Technology Value Matrix. The Alteryx One platform delivers analytics automation by enabling organizations to prepare data, define business logic, and operationalize AI-driven workflows across the enterprise. Available in Starter, Professional, and Enterprise editions, the platform aligns to organizational maturity and scale while supporting a broad range of use cases, from analyst-driven insights to automated, production-grade processes.

Alteryx emphasizes accessibility for business users while maintaining the depth required by technical teams. Through its visual workflow interface, Alteryx Designer enables users to build repeatable, auditable analytics processes without code, while Designer Cloud extends these capabilities into cloud-native environments. AI capabilities, delivered through AI Builder and conversational experiences such as Ask Alteryx, are embedded directly into workflows, helping users build, understand, and automate processes using both predictive and generative AI.

A key differentiator is Alteryx's focus on operationalizing analytics. Rather than stopping at insights, the platform enables organizations to translate analytics into governed, repeatable actions. Enterprise features such as role-based access controls, workflow versioning, and auditability ensure that data, logic, and AI outputs remain trusted and aligned with IT governance. This approach allows organizations to scale analytics and AI adoption while maintaining control, transparency, and consistency.

Recent product updates and enhancements include:

► Alteryx One

New tiered platform packaging (Starter, Professional, Enterprise) unifying previously separate products into a single scalable offering aligned to customer adoption maturity.

Alteryx's new tiered platform packaging unifies previously separate products into a single consumption-based model, making enterprise analytics accessible at every organizational level.

Alteryx Copilot Gen 2 expands AI-assisted workflow construction across all core editions, lowering the barrier for citizen data scientists to build production-grade analytics.

▶ **Consumption-Based Pricing**

Alteryx Server automation capabilities shifted from core-based to automation-run pricing, improving scalability and aligning cost with actual usage delivered.

▶ **AI Builder Bundled**

Intelligence Suite and generative AI capabilities combined into AI Builder and included by default in Professional and Enterprise editions, reducing adoption friction for advanced analytics and GenAI use cases.

▶ **Alteryx Copilot Gen 2**

Expanded AI assistant integrated across core editions for workflow construction, explanation, and acceleration using generative AI, included by default in Professional and Enterprise.

▶ **Cloud-First Expansion**

Continued enhancements to Designer Cloud, Cloud Execution, and Plans with maintained on-premises and hybrid support, giving customers clearer paths to cloud modernization.

▶ **Simplified User Roles**

Formalized Full Creator, Basic Creator, and Viewer roles to better reflect real usage patterns and simplify organizational scaling of analytics.

Alteryx's shift to consumption-based pricing for Server automation reflects a broader market trend toward usage-aligned licensing that reduces upfront cost barriers for enterprise adoption.

Amazon

Amazon Web Services (AWS) is recognized as a leader in this year's DSML Platform Technology Value Matrix. Its machine learning and AI services form a comprehensive ecosystem for developing, deploying, and managing models across the full analytics lifecycle. SageMaker functions as the central development environment, with SageMaker Studio providing integrated notebook experiences, automated experiment tracking, and collaborative development features. SageMaker Data Wrangler simplifies data preparation with over 300 built-in transformations, while SageMaker Feature Store provides centralized feature management for consistent model training and inference.

AWS addresses diverse analytical requirements through Amazon EMR for large-scale data processing, AWS Glue for ETL operations, and Amazon Athena for SQL-based analytics, all integrating with SageMaker into a unified data science pipeline. SageMaker Autopilot automates algorithm selection and hyperparameter tuning, while

Amazon Bedrock AgentCore provides composable agent infrastructure supporting multi-hour async workloads, short- and long-term memory, and production-grade identity and guardrails.

SageMaker JumpStart provides pre-trained models and deployment frameworks. Amazon Bedrock extends the platform as a managed service for accessing, customizing, and deploying foundation models, with Bedrock Guardrails providing content filtering and governance controls. SageMaker Clarify, Model Monitor, and fine-grained access controls support bias detection, continuous quality assessment, and regulatory compliance.

Recent product updates and enhancements include:

► **SageMaker Unified Studio GA**

A unified data and AI development environment combining EMR, Glue, Athena, Redshift, Bedrock, and SageMaker AI, with one-click onboarding, serverless notebooks, and a built-in AI agent.

► **Amazon Bedrock AgentCore**

Composable agent infrastructure with runtime supporting multi-hour async workloads, short- and long-term memory, and identity integration with Cognito, Entra ID, and Okta. Multi-agent collaboration enables supervisor-coordinated workflows.

► **SageMaker AI Serverless Customization**

Serverless model customization with RLVR and RLAIIF on a pay-per-token model, bidirectional streaming for real-time voice agents, and rolling deployment updates with automatic rollback.

Databricks

Databricks is positioned as a leader in this year's DSML Platform Technology Value Matrix. The Lakehouse Platform unifies data engineering, data science, machine learning, and analytics in a single architecture, enabling organizations to build and operationalize AI systems while maintaining governance and security at scale. Delta Lake technology underpins data processing for both batch and streaming workloads, while MLflow manages the machine learning lifecycle from experimentation through deployment with standardized reproducible workflows.

The platform supports multiple development approaches through notebooks in SQL, Python, R, and Scala, accommodating data engineers building automated pipelines, data scientists implementing custom models, and business analysts using natural language query interfaces. AutoML automates algorithm selection, feature engineering, and hyperparameter optimization with transparency into the generated pipeline, while the Feature Store maintains consistent

SageMaker Unified Studio combines EMR, Glue, Athena, Redshift, and Bedrock into a single development environment, eliminating the fragmentation that previously slowed enterprise ML teams.

Databricks' Mosaic AI Agent Framework delivers production-scale agent building and evaluation with Unity Catalog governance and tool-use tracing built in.

feature definitions between training and inference. Unity Catalog provides centralized governance for data, AI, and analytics assets, including discovery, lineage tracking, and access controls. Mosaic AI extends the platform with capabilities for building, customizing, and deploying AI models at production scale.

Recent product updates and enhancements include:

► **Mosaic AI Agent Framework**

A production-scale framework for building, deploying, and evaluating AI agents with Unity Catalog governance, MLflow tracing, Vector Search for RAG, and support for any orchestration library.

► **AI/BI Genie Agentic Analytics**

Agentic experience with Agent mode (Beta) featuring parallel queries, thinking traces, natural-language dashboard authoring, and an assistant for automating multi-step analytical workflows.

► **Genie Research Agent**

Enhanced the Research Agent (Beta) to support hypothesis-driven analysis using multi-step reasoning. The agent creates research plans, executes multiple SQL queries, and iteratively refines its approach based on results to produce comprehensive analytical summaries for complex data questions.

AI/BI Genie's Agent mode enables parallel queries with thinking traces and natural language explanations, transforming business intelligence from static dashboards to interactive analytical conversations.

Dataiku

Dataiku is recognized as a leader in this year's DSML Platform Technology Value Matrix. Its AI platform covers the development, deployment, and management of data science projects within a single end-to-end environment, enabling organizations to build, deploy, monitor, and govern analytics, machine learning models, and AI agents while reducing operational complexity. The platform supports both visual workflows and code-based development across user skill levels. AutoML accelerates model creation, while advanced users can implement custom models using Python, R, and SQL with integrations to TensorFlow, PyTorch, and scikit-learn.

Dataiku serves both technical practitioners and business decision-makers. Data scientists and engineers benefit from integrated experimentation and deployment tooling, while business leaders gain oversight through dashboards, monitoring, and governance capabilities. Role-based access controls, model monitoring, automated documentation, and audit trails support compliance and

Dataiku's Agent Hub enables multi-agent orchestration across Snowflake Cortex, Databricks, and Google Vertex AI, with Agent Review providing continuous quality validation.

transparency, while explainability and bias detection features maintain accountability in production models. Performance tracking and drift detection ensure ongoing reliability. Generative AI capabilities are embedded throughout for data augmentation, reporting, and workflow optimization within governed boundaries, with multi-agent orchestration now a core part of the platform architecture.

Recent product updates and enhancements include:

► **Dataiku 14.x Platform**

Major release cycle introducing Agent Hub for multi-agent orchestration, third-party and custom Python agent tools, unified LLM guardrails, LLM cost blocking and alerting, reranking for RAG and agentic retrieval, and semantic models for agents.

► **Agent Review**

Launched in Dataiku 14.4, enabling AI builders and subject matter experts to define shared quality standards, build repeatable tests, and validate agent behavior continuously with optional SME grading.

► **External Agent Orchestration**

Dataiku now manages and composes agents across Snowflake Cortex, Databricks, Google Vertex AI, and AWS Bedrock as governed first-class agents. Approved agents can be deployed directly into Slack.

Dataiku 14.x introduced Agent Hub for multi-agent orchestration and third-party plugin support, positioning the platform as an agent composition layer across the enterprise AI stack.

SAS

SAS is recognized as a leader in this year's DSML Platform Technology Value Matrix. The SAS Viya platform provides an integrated analytics environment spanning data preparation, visualization, statistical analysis, and machine learning across the full data science lifecycle. Designed for enterprise deployment, the platform supports both visual interfaces and programmatic development through SAS, Python, R, and Java, accommodating users with different technical backgrounds. For machine learning, SAS Viya combines automated model generation with customization options for advanced practitioners, supporting experimentation, development, and deployment within a governed environment.

SAS Viya serves technical specialists and business stakeholders alike. Data scientists access integrated tools for model development while business leaders interpret results through dashboards and visualizations. Enterprise governance capabilities include access

SAS Viya Intelligent Decisioning enables configurable human-AI oversight ratios for agent deployments, allowing organizations to calibrate autonomy levels to their risk tolerance.

controls, lineage tracking, model monitoring with drift detection, and explainability features supporting regulatory compliance. Expanded generative AI capabilities assist with code generation, data preparation, and workflow development within controlled environments, with intelligent decisioning and synthetic data generation extending the platform's AI capabilities into new enterprise use cases.

Recent product updates and enhancements include:

► **SAS Viya Copilot GA**

Generally available in the 2025.12 release, supporting code assistance, model pipeline development in Model Studio, and Model Manager, powered by GPT-4.1 and Phi-3 models on Azure AI Foundry.

► **SAS Viya Intelligent Decisioning**

Enables building and deploying intelligent AI agents with configurable human-AI oversight ratios, allowing organizations to calibrate automation autonomy against task complexity and operational risk.

► **SAS Data Maker**

A secure synthetic data generator addressing data privacy and scarcity challenges, accelerated by SAS's acquisition of Hazy's principal software assets, generally available in Q3 2025.

SAS Data Maker provides secure synthetic data generation to address privacy and scarcity challenges, accelerated by the company's 50 years of statistical expertise in data modeling.

Teradata

Teradata is positioned as a leader in this year's DSML Platform Technology Value Matrix. VantageCloud provides an integrated environment for data science and machine learning, enabling practitioners to develop and deploy models through a unified architecture optimized for complex analytical workloads. Its parallel processing architecture reduces data movement by executing operations within the database, allowing data scientists to work with complete datasets. ClearScape Analytics serves as the analytical core, delivering statistical and mathematical functions through Python, R, and SQL interfaces, including Jupyter based workflows.

The platform supports machine learning development through in-database functions as well as open-source frameworks and tools, spanning traditional statistical models to deep learning, with automated feature engineering, hyperparameter optimization, and support for techniques such as gradient boosting, random forests, and

Teradata's Enterprise AgentStack and open-source MCP Server standardize agent-to-model communication, enabling scalable agent deployment governed by enterprise security frameworks.

neural networks. An Enterprise Feature Store enables reuse of ML features across models and projects, the Enterprise Vector Store supports embeddings and hybrid search across structured and unstructured data. Teradata serves data engineers, data scientists, business analysts, and executives through differentiated interfaces, with ModelOps supporting model lifecycle management and performance monitoring in production. Teradata is also extending this foundation into agentic AI with Enterprise AgentStack and related capabilities for building, deploying, and governing enterprise AI agents grounded in trusted data, context, and analytics. Recent partner announcements, including Unstructured and NVIDIA, expand support for unstructured data ingestion and high-performance AI workloads.

Recent product updates and enhancements include:

▶ **Teradata Data Analyst Agent**

An enterprise-grade AI analytics agent available on Google Cloud Marketplace, enabling advanced agentic AI integration directly into customer cloud environments.

▶ **Enterprise AgentStack**

A unified platform to build, deploy, and govern enterprise AI agents at scale, including modular components such as AgentBuilder, AgentOps, and Enterprise MCP, for structured agent lifecycle management.

▶ **AgentBuilder with Expert Agents**

A no-code and pro-code environment for creating intelligent agents using expert agent templates and guided workflows for enterprise automation use cases.

▶ **Teradata AI Microservices**

A modular set of containerized services within the Teradata AI Factory enabling scalable AI pipelines, vector functions, and integration with third-party AI tools.

▶ **Enterprise MCP Server**

An open-source Model Context Protocol server standardizing AI agent-to-model communication for scalable agentic architectures across enterprise environments.

▶ **Enterprise Vector Store with Hybrid Search**

Teradata AI Microservices offer containerized, modular AI services within the AI Factory framework, enabling organizations to scale AI pipelines without rearchitecting their data infrastructure.

Teradata's Enterprise Vector Store with native hybrid search supports semantic similarity within the existing data warehouse, eliminating the need for separate vector database infrastructure.

In-database vector storage and native hybrid search supporting semantic similarity, RAG, and AI agent use cases at enterprise scale.

Experts

Experts in the DSML Platform Technology Value Matrix include Google, MathWorks, Microsoft, Oracle, and Palantir.

Google

Google is recognized as an expert in this year's DSML Platform Technology Value Matrix. The Vertex AI platform integrates machine learning development, deployment, and operations within a unified architecture connecting data processing, model development, and production management. Vertex AI Feature Store centralizes feature definitions with batch and real-time access patterns, while Vertex AI Pipelines enables reproducible data processing through containerized components. For model development, Vertex AI Workbench provides interactive notebook environments with pre-configured dependencies, accommodating both AutoML for automated model creation and Custom Training for specialized requirements.

The Gemini foundation model framework is central to Google's DSML strategy, integrated throughout Vertex AI with multimodal capabilities across text, code, images, and audio, accessible through standardized APIs. Model Garden provides specialized Gemini tuning for domain-specific applications, while Vertex AI Tuning offers optimization workflows preserving core model capabilities. Vertex Explainable AI, Vertex AI Model Monitoring, and Vertex AI Safety Tooling complete a governance framework spanning attribution methods, drift detection, and content filtering. Google's agentic AI infrastructure, anchored by the Agent Development Kit, Agent Engine, and A2A protocol, has become a defining element of the platform's 2025 direction.

Recent product updates and enhancements include:

- ▶ **Agent Development Kit, Agent Engine, and A2A Protocol**
Open-source ADK for building multi-agent systems in Python, plus a managed Agent Engine runtime on Vertex AI with memory management and evaluation tools. The Agent2Agent (A2A) open protocol enables cross-platform agent interoperability backed by 50+ partners.
- ▶ **Agentspace and Agent Designer**

Google's Agent Development Kit, Agent Engine, and Agent2Agent protocol form a comprehensive framework with over 50 partners committed to cross-platform agent communication.

Agentspace provides unified enterprise search with pre-built agents for Deep Research, NotebookLM, and custom agent deployment, bridging consumer AI capabilities with enterprise governance requirements.

Agentspace provides unified enterprise search with pre-built Google agents (Deep Research, Idea Generation, NotebookLM for Enterprise) and a no-code Agent Designer allowing non-technical employees to create custom agents.

► **Vertex AI Infrastructure Expansion**

Ironwood TPU and A4X Max GPU instances (NVIDIA GB300 NVL72) launched. Vertex AI became the only platform with generative media models across video (Veo), image (Imagen 3), speech, and music (Lyria).

MathWorks' MATLAB Copilot brings generative AI code assistance to engineering and scientific computing, bridging domain-specific modeling with modern AI workflows.

MathWorks

MathWorks is positioned as an expert in this year's DSML Platform Technology Value Matrix. The MATLAB and Simulink platforms provide a technical environment for data science, machine learning, and AI development with particular emphasis on integrating AI with simulation for systems-level applications, including model-based design, embedded AI, and safety-critical implementations. MATLAB delivers a unified programming environment with extensive mathematical and statistical functions, while specialized toolboxes, Statistics and Machine Learning Toolbox, Deep Learning Toolbox, and Reinforcement Learning Toolbox, span traditional statistical models through advanced neural network architectures.

For data preparation, MATLAB provides functions for missing value handling, outlier detection, feature engineering, and dimensionality reduction. The platform supports deployment across diverse environments through MATLAB Production Server, standalone applications via MATLAB Compiler, and embedded implementations through code generation for FPGAs, microcontrollers, and embedded GPUs and CPUs. Simulink integration enables simulation-driven ML development where training and validation leverage synthetic data, particularly valuable where physical data collection is impractical. Governance and traceability mechanisms track datasets, algorithms, and parameters throughout the development process.

MATLAB's R2025a release added SAM-based ground truth labeling and incremental k-means for streaming data, extending the platform's relevance to real-time edge AI applications.

Recent product updates and enhancements include:

► **MATLAB Copilot (R2025b)**

A generative AI assistant integrated into MATLAB with code autocompletion, natural-language code generation, error explanation, test generation, and documentation-sourced Q&A.

► **R2025a AI and Deployment Enhancements**

SAM-based ground truth labeling in Video Labeler, incremental k-means for streaming data, improved PyTorch and TensorFlow model import performance, and new deployment targets for Qualcomm and Infineon NPU hardware.

► **EDGE AI Foundation Membership**

MathWorks joined the EDGE AI Foundation to advance energy-efficient embedded AI, formalizing an end-to-end workflow for training, simulation, compression, and code generation to CPUs, GPUs, FPGAs, and NPUs from a single Simulink model.

Microsoft Foundry launches with 1,400+ agentic integrations and visual multi-agent workflows, establishing a unified portal for enterprise AI agent development.

Microsoft

Microsoft is positioned as an expert in this year's DSML Platform Technology Value Matrix. Azure Machine Learning delivers an integrated environment for data science and machine learning, enabling organizations to develop, deploy, and manage AI solutions with enterprise governance across the full analytics lifecycle. The service supports both automated machine learning workflows and custom model development, with interactive notebooks in Python, R, and Scala, experiment tracking for methodical comparison, and integration with scikit-learn, PyTorch, TensorFlow, and XGBoost. Azure ML Pipelines allow organizations to create reusable workflows that standardize data preparation, model training, and evaluation processes.

For model deployment, Azure provides real-time endpoints, batch inference, and edge deployment options. Model monitoring tracks performance metrics, data drift, and resource utilization, while the model registry supports versioning and lifecycle management integrated with Azure Kubernetes Service. The Microsoft Foundry platform, rebranded and expanded from Azure AI Studio, serves as the central hub for agentic AI development, providing foundation model access, multi-agent visual workflows, and a broad set of agentic integrations across the Azure AI ecosystem. Role-based access controls and fairness assessment capabilities support responsible AI governance throughout the platform.

Agent 365 introduces a new control plane to manage and secure agents across Office applications with enterprise-grade audit trails, usage analytics, and DLP policy enforcement.

Recent product updates and enhancements include:

► **Microsoft Foundry Launch**

A unified portal with 1,400+ agentic integrations, multi-agent visual workflows, and a new Azure AI SDK unifying agents, inference, evaluations, and memory. GPT-5.2 and Codex Max reached general availability.

► **Foundry Agent Service**

Managed agent hosting, MCP server (preview), Agent2Agent tool (preview), managed long-term memory, and one-click publishing to Microsoft 365 Copilot and Teams. Foundry IQ provides enterprise knowledge grounding without custom RAG pipelines.

► **Agent 365 and Copilot Expansion**

Agent 365 introduced as a new control plane to manage and secure agents across Office applications. Power BI Copilot expanded to all Fabric SKUs, with Copilots added for Dynamics 365 Sales, Supply Chain, Finance, and Customer Service.

Oracle's Select AI Agent introduces an in-database framework for building autonomous agentic workflows using PL/SQL, integrating agents directly into the database layer.

Oracle

Oracle is positioned as an expert in this year's DSML Platform Technology Value Matrix. The Oracle Cloud Infrastructure (OCI) Data Science platform delivers an end-to-end environment for developing, deploying, and managing AI and machine learning solutions with enterprise-grade governance. The platform integrates data management with analytics capabilities, supporting connectivity to Oracle databases, object storage, and third-party repositories, with distributed computing frameworks for large-scale dataset processing. OCI Data Science provides algorithm implementations across supervised learning, unsupervised techniques, and deep neural networks.

Oracle's AutoML incorporates automated feature selection, algorithm evaluation, and hyperparameter optimization, while accelerators for computer vision, natural language processing, and recommendation systems integrate pre-trained components with domain-specific customization. JupyterLab notebooks support exploratory analysis with Python and specialized frameworks in collaborative workspaces. Governance features include lineage tracking, audit mechanisms, drift detection, and automated retraining workflows. OCI Generative AI service and Data Science AI Quick Actions extend the platform with foundation model access, and Select AI Agent introduces in-database agentic workflow capabilities, a notable architectural evolution for Oracle's offering.

Select AI Agent enables autonomous agentic workflows using PL/SQL within the database, reducing the architectural complexity of deploying AI agents in data-intensive enterprise environments.

Recent product updates and enhancements include:

► **AI Quick Actions v2.0**

Support for OpenAI endpoint model deployment Oracle, multimodel serving and model groups, tool calling via vLLM, and a no-code interface for deploying, fine-tuning, and evaluating

foundation models including Llama 4, Granite, and open-weight models from OpenAI and others.

► **Select AI Agent**

An in-database framework for building, deploying, and managing autonomous agentic AI workflows using PL/SQL and Python, with built-in row, column, and cell-level data privacy protections.

► **ML Applications**

A new capability that standardizes the packaging and deployment of AI/ML functionality, enabling organizations to build, deploy, and operate machine learning as a service. Oracle Cloud Infrastructure It provides end-to-end lifecycle management from development through production with fleet management at scale.

► **Distributed Jobs and MLOps Enhancements**

Distributed jobs offer on-demand scalability and cost efficiency Oracle for training workloads. Additional MLOps updates include ML Pipeline scheduling, Data Flow integration in ML Pipelines, and burstable VM support for jobs, notebooks, and pipelines.

Palantir

Palantir is positioned as an expert in this year's DSML Platform Technology Value Matrix. The Palantir Foundry platform provides a data science and machine learning environment integrating data analysis, model development, and operational deployment within a unified enterprise-scale architecture. Foundry's ontology layer establishes semantic relationships between data entities, creating a contextual data foundation that maintains lineage across analytical workflows. For machine learning development, Foundry supports both visual modeling interfaces and code-based environments in Python, R, and SQL, accommodating users with varying technical capabilities.

Palantir AIP (Artificial Intelligence Platform) connects large language models with enterprise data systems within appropriate security controls, enabling context-aware applications drawing on both public and proprietary information sources. AIP Agent Studio serves as the primary environment for designing, governing, and deploying enterprise AI agents with integration into Foundry's ontology for contextual awareness. Governance capabilities span granular permission management, audit trails, documentation features, and model monitoring frameworks that track performance and manage versioning for production reliability.

Palantir's Ontology MCP enables AI agents to autonomously interact with enterprise data objects, bridging the gap between external AI models and proprietary operational data.

Palantir's agent monitoring provides near real-time metrics on success rates, latency, and cost per action, enabling organizations to optimize agent performance at enterprise scale.

Recent product updates and enhancements include:

► **AIP Agent Studio Enhancements**

Agents can now be published as Functions for use in AIP Evals, Automate, and Code Repositories. Native tool calling with parallel tool calls shipped in August 2025, improving agent execution speed.

► **Palantir MCP and Ontology MCP**

Palantir MCP enables AI IDEs and agents to autonomously design, build, and review applications within the platform. Ontology MCP exposes Developer Console applications as discoverable tools for any MCP-compatible agent framework.

► **Agent Monitoring and Observability**

New monitoring views for functions and actions with near real-time metrics (success and failure rates, P95 duration) and multi-channel alert notifications for proactive reliability management.

Siemens' Agent Studio with MCP integration enables agents to directly query, reason over, and interact with Graph Studio knowledge graphs for complex analytical workflows.

Accelerators

Accelerators in the DSML Platform Technology Value Matrix include Siemens (Altair), DataRobot, H2O.ai, Hex Technologies, and TrueFoundry.

Siemens (Altair)

Siemens is positioned as an accelerator in this year's DSML Platform Technology Value Matrix. The Altair RapidMiner platform delivers an integrated environment for data science and machine learning, supporting the complete analytical lifecycle. Data engineering capabilities are delivered through standardized interfaces and knowledge graphs that integrate data into a semantic layer, with visual workflow-based transformation complemented by Python, R, SAS, and SQL for specialized requirements. For machine learning, the platform provides automated model generation and customizable development environments supporting diverse algorithms and methodologies.

Data scientists benefit from collaborative notebook environments, version control, and experiment tracking, while business analysts utilize visual interfaces that abstract technical complexity. Decision-makers access interactive dashboards translating analytical outputs into actionable insights. The governance framework includes model documentation, version control, lineage tracking, and explainability

Agent Studio combines LLMs, ML models, and enterprise data sources into orchestrated agentic workflows, with MCP integration enabling direct interaction with Graph Studio knowledge graphs.

through feature importance visualization. Siemens has expanded its AI capabilities into agentic workflow construction and knowledge graph integration via Model Context Protocol.

Recent product updates and enhancements include:

► **Agent Studio**

A new environment for building and orchestrating agentic workflows combining LLMs, ML, and enterprise data, with multi-agent collaboration capabilities for reasoning, retrieval, and task automation.

► **Graph Studio with MCP Integration**

Agents interact directly with Graph Studio via Model Context Protocol to query, reason, and make decisions. LLM-powered copilots support metadata management, ontology creation, and attribute-based access controls.

DataRobot

DataRobot is positioned as an accelerator in this year's DSML Platform Technology Value Matrix. The DataRobot AI Platform delivers an end-to-end enterprise environment for developing, deploying, and managing machine learning models at scale. The platform's automated machine learning capabilities accelerate model development through automated feature engineering and a model blueprint system that evaluates multiple algorithm combinations for specific business problems. Technical users can extend these capabilities through Python, R, and other programming interfaces for specialized methodology integration.

DataRobot serves diverse organizational roles, data scientists accelerate experimentation while maintaining methodological rigor, business analysts implement predictive solutions through guided interfaces, and IT teams use deployment frameworks with monitoring capabilities. The governance architecture incorporates model documentation, compliance tracking, drift detection, explainability, and role-based access management. MLOps capabilities automate containerization, scaling, and monitoring while supporting measurement of business impact for production deployments. A full agentic AI platform now extends the offering beyond traditional AutoML.

Recent product updates and enhancements include:

DataRobot's agentic platform supports bring-your-own agent frameworks including LangGraph, LlamaIndex, and CrewAI, with standardized MCP connectivity for enterprise

DataRobot's MCP Server support with clone-and-deploy architecture enables standardized connectivity between agents and enterprise data sources without custom integration work.

► **Agentic AI Platform**

A full end-to-end agentic platform with support for BYO agent frameworks (LangGraph, LlamaIndex, CrewAI), an agentic playground, OTEL-compliant tracing, and one-click NVIDIA NIM deployment in air-gapped environments.

► **MCP Server Support and Talk to My Docs**

Standardized connectivity via Model Context Protocol with a clone-and-deploy MCP server template. Talk to My Docs provides governed agentic RAG across Google Drive, Box, and local files.

► **Expanded LLM Gateway**

Added Cerebras and Together AI alongside existing providers, with access to Gemma, Mistral, and Qwen models through a single governed interface. CLI tooling and agentic app starter templates streamline local-to-production deployment.

H2O.ai's h2oGPTe achieved the top ranking on the GAIA leaderboard in March 2025, validating its agentic AI capabilities against the industry's most rigorous benchmark.

H2O.ai

H2O.ai is positioned as an accelerator in this year's DSML Platform Technology Value Matrix. The H2O AI Cloud platform delivers a unified environment for data science and machine learning enabling organizations to build, deploy, and operate AI solutions across diverse business applications. The platform provides visual workflows for business users through H2O Flow and code-based environments for technical specialists via H2O-3 and Python integration. H2O AutoML enables automated model generation while experienced practitioners implement custom algorithms through the platform's extensible framework. Driverless AI supports experimentation, model development, and deployment with built-in feature engineering capabilities.

The H2O AI Feature Store, commercially launched with AT&T, delivers sub-millisecond feature serving for real-time AI applications across financial services and telecommunications.

H2O.ai serves multiple stakeholder groups, with H2O Wave dashboard applications providing executive-level visibility into AI initiative status and business impact. The platform incorporates enterprise governance through H2O Model Documentation, MLOps monitoring with drift detection, and Machine Learning Interpretability (MLI) functions supporting transparency in regulated applications. H2O LLM Studio and h2oGPTe extend the platform with generative AI and agentic capabilities, including a top-ranked performance on the GAIA leaderboard, within a governed enterprise framework.

Recent product updates and enhancements include:

► **h2oGPTe Agentic Capabilities**

Achieved the #1 ranking on the GAIA leaderboard (March 2025). Added Tool Builder Agent, human-in-the-loop planning, automatic model routing, agent monitoring, and Final Answer Mode controls.

► **H2O AI Feature Store Commercial Launch**

Joint commercialization with AT&T of the H2O AI Feature Store, providing sub-millisecond latency for production transactions, automatic feature recommendations, drift and bias detection, and FeatureRank scoring, now available as a standalone enterprise product.

► **Enterprise h2oGPTe Expansions**

Added role-based access controls, OpenAPI-compliant REST API, Auto-Sync Collection Connector, SharePoint OAuth2 authentication, RAG query filters, and MCP enhancements. Secure Amazon Bedrock integration enables private VPC model access.

Hex Technologies

Hex Technologies is positioned as an accelerator in this year's DSML Platform Technology Value Matrix. Its collaborative data platform unifies data science capabilities with accessible interfaces, enabling organizations to transform analytical workflows into interactive applications. The Hex platform provides notebook-based workspaces combining SQL, Python, and visualization capabilities, with transformation tools for data cleaning, shaping, and enrichment through both code-based and visual methods. For machine learning, Hex supports scikit-learn, PyTorch, and TensorFlow with infrastructure for systematic experimentation, hyperparameter optimization, and metric tracking.

Hex serves multiple user types, data scientists with ML tooling and deployment workflows, ML engineers with automated retraining and model serving capabilities, analysts through guided interfaces, and business stakeholders through intuitive applications without requiring technical expertise. Hex Magic, the platform's integrated AI assistant, leverages large language models to assist with code development, query optimization, visualization design, and narrative generation throughout the workflow. The platform's Fall 2025 release represented its most expansive product update to date, introducing conversational analytics and an agentic Notebook experience.

Recent product updates and enhancements include:

Hex's Threads and Notebook Agent represent a major evolution in collaborative analytics, combining conversational self-serve interfaces with automated model-building agents.

Hex's AI Modeling Agent inspects existing projects and synthesizes code into semantic models with diff-based review, automating one of the most time-consuming steps in analytics workflows.

► **Threads and Notebook Agent**

Hex's largest release shipped Threads (a conversational self-serve analytics interface), an upgraded agentic Notebook, and semantic model support grounded in curated context.

► **AI Modeling Agent**

A new agent that inspects existing Hex projects and synthesizes code into semantic models, with a diff view for reviewing proposed changes and cross-model calculation support.

TrueFoundry

TrueFoundry is positioned as an accelerator in this year's DSML Platform Technology Value Matrix. The TrueFoundry platform provides MLOps infrastructure designed to streamline the deployment, management, and scaling of machine learning models and AI applications. Its architecture emphasizes operational efficiency and infrastructure abstraction through containerization and Kubernetes orchestration, enabling consistent ML operations regardless of the underlying technology stack. The platform provides tools for model deployment across diverse environments, automated monitoring and observability, and centralized management of ML workloads.

TrueFoundry serves technical practitioners with deployment frameworks and monitoring tools for model operationalization, as well as business leaders tracking application performance through centralized dashboards. Governance features include model versioning, performance tracking, and audit trails, with role-based access controls and configurable permissions. The AI Gateway for production LLMs, featuring intelligent model routing, request caching, and budget governance, extends the platform's capabilities to large-scale language model deployments. Enterprise guardrails and MCP governance now support governed agentic system deployments in regulated environments.

Recent product updates and enhancements include:

► **Advanced AI Gateway for Production LLMs**

Virtual Models with intelligent routing across multiple providers, latency-based and weighted routing policies, automatic retries with model fallback, request caching (exact and semantic), and enhanced budget governance with weekly limits and alerts.

► **Enterprise Guardrails and MCP Governance**

TrueFoundry's Enterprise Guardrails with MCP governance provide secret detection, SQL sanitization, and PII handling at the gateway layer, securing agent interactions at scale.

TrueFoundry's Virtual Models with intelligent routing across multiple providers enable latency-based failover and cost optimization, giving organizations provider independence for LLM deployments.

Managed Guardrails including Secret Detection, SQL Sanitization, PII detection, Content Moderation, and Prompt Injection protection. MCP Guardrails at the Gateway layer enable policy enforcement across MCP servers and models for governed agentic deployments.

► **Infrastructure Automation and Runtime Optimization**

Enhanced autoscaling and spot/on-demand orchestration, fractional GPU support improvements, autopilot-driven infrastructure recommendations, and Git tag-based deployments for improved release control.

► **Identity and Access Controls**

SCIM support for automated user provisioning, identity provider revamp with expanded SSO integration, custom roles, and expanded RBAC controls supporting multi-team AI adoption at scale.

► **Deep Observability and Data Governance**

Data Routing Rules for control-plane or customer-managed log storage supporting data residency requirements, expanded OpenTelemetry support, and token-level cost and usage analytics.

Core Providers

Core Providers in the DSML Platform Technology Value Matrix include Cloudera, Domino Data Lab, IBM, and Samsung.

Cloudera

Cloudera is positioned as a core provider in this year's DSML Platform Technology Value Matrix. The Cloudera Data Platform (CDP) delivers an enterprise data cloud unifying data engineering, data warehousing, machine learning, and analytics in a single integrated ecosystem across hybrid and multi-cloud environments. Cloudera Data Engineering facilitates scalable data processing and transformation via Apache Spark with automated pipeline orchestration, while the SDX (Shared Data Experience) layer maintains consistent security policies, governance controls, and metadata management across workloads.

Cloudera Machine Learning (CML) provides a collaborative workspace for data scientists to explore data, build models, and deploy AI applications, supporting Python, R, and Scala alongside visual low-code interfaces. The platform serves data engineers, data scientists, business analysts, and executives through differentiated interfaces.

Cloudera's on-premises AI inference with embedded NVIDIA NIM microservices addresses a critical gap for organizations requiring sovereign AI deployment with full data control.

Cloudera's Data Warehouse with Trino on-premises extends centralized security and governance across the full data estate, bridging the gap between cloud-native and on-premises analytics.

The governance framework uses Atlas for metadata management and lineage tracking, and Ranger for fine-grained access controls and audit capabilities to support regulatory compliance. On-premises AI inference capabilities have been extended to data centers in 2026 as a meaningful step toward enterprise AI deployment flexibility.

Recent product updates and enhancements include:

► **Cloudera AI Inference On-Premises**

One of the first AI inference services with embedded NVIDIA NIM microservices, now available in data centers powered by Blackwell GPUs, Dynamo-Triton Inference Server, and Nemotron open models.

► **Cloudera Data Warehouse with Trino On-Premises**

Extends centralized security, governance, and observability across the full data estate with integrated AI-powered analytics and the Iceberg REST Catalog.

► **AI Studios and Data Visualization Enhancements**

AI Studios (May 2025 technical preview) provides low-code tools for generative AI development. Data Visualization added AI annotation for auto-generating chart summaries, AI query logging and traceability, and resilient AI features.

Domino Data Lab

Domino Data Lab is positioned as a core provider in this year's DSML Platform Technology Value Matrix. The Domino enterprise AI platform provides a unified environment for data science and machine learning development, centralizing data science operations while providing flexibility for diverse modeling approaches and technical requirements. The platform integrates data connections through standardized interfaces supporting visual workflows and code-based development with Python, R, and SAS. Template-based model development is complemented by customization options that balance productivity with technical precision. Containerization and Kubernetes orchestration support consistent functionality across both on-premises and cloud deployments.

Domino serves technical practitioners and business stakeholders, with data scientists accessing integrated experimentation and deployment tools while business leaders monitor project progress through dashboards. Governance features include access controls, model versioning, audit capabilities, and monitoring with drift detection for

Domino's Agentic Development Lifecycle Platform is the first fully governed end-to-end platform for agentic AI, featuring a universal agent gateway and compliance monitoring.

Domino's AI Factory capabilities standardize model development workflows and streamline handoffs between data science and engineering teams, reducing time-to-production for enterprise models.

production reliability. The AI Gateway simplifies and secures access to external LLMs, while Domino Governance automates compliance tracing within the data scientist's workflow. The Winter 2026 release introduced a fully governed Agentic Development Lifecycle platform, marking a meaningful expansion in Domino's AI capabilities.

Recent product updates and enhancements include:

► **Agentic Development Lifecycle Platform**

The first fully governed end-to-end platform for agentic AI, featuring a universal tracing SDK, structured side-by-side evaluation, production-ready deployment with autoscaling, continuous monitoring, and governed LLM hosting.

► **AI Factory Capabilities**

Structured workflows standardizing model development and streamlining handoffs between data science, IT, and risk teams, automating the experimentation-to-production pipeline. Also launched Zero-to-AI consulting service.

► **Autoscaling Apps and App Discovery**

Domino Apps now elastically scale compute to meet demand, with a new discovery experience for business users and infrastructure cost reductions of up to 60% in some deployments.

IBM

IBM is positioned as a core provider in this year's DSML Platform Technology Value Matrix. The IBM Watson Studio platform provides an integrated environment for data science and machine learning, enabling organizations to build, train, and manage AI models at scale. Watson Studio supports multiple development approaches through visual interfaces and code-based workflows in Python, R, and Scala, with data management tools for connecting, preparing, and transforming data from diverse sources. Automated model building is combined with customization options for practitioners of varying technical backgrounds, while experiment tracking and version control support iterative development.

The watsonx platform extends IBM's capabilities as the foundation for enterprise-grade AI implementation, offering pre-trained foundation models with domain specializations for financial services, healthcare, and other industries. Watson OpenScale provides monitoring for deployed models with bias detection and explainability support. The watsonx ecosystem enables prompt engineering, fine-tuning, and

IBM's watsonx Orchestrate now includes agent observability and governance capabilities, providing visibility into agent reasoning chains for enterprise compliance requirements.

IBM Data Product Hub and Data Gate for watsonx introduce structured data product management, enabling organizations to treat curated datasets as discoverable, governed enterprise assets.

RAG workflows connecting foundation models with enterprise knowledge sources. The watsonx AI Model Gateway provides unified access to models from OpenAI, Anthropic, and NVIDIA across diverse deployment environments, and agentic capabilities have been added through watsonx Orchestrate.

Recent product updates and enhancements include:

► **Platform Expansion Announcements**

At IBM Think 2025, IBM announced new data and automation capabilities across the watsonx platform, including broadened collaborations with hyperscalers and model providers and enhanced support for open-source foundation models through the Granite family and partnerships with providers such as Meta.

► **watsonx Enhancements**

New features streamline the AI development lifecycle and provide greater flexibility and control for data scientists, AI builders, and developers, including capabilities announced at IBM Think 2025.

► **Data Products and Integration**

IBM introduced IBM Data Product Hub and Data Gate for watsonx, along with new integration and intelligence capabilities in watsonx.data to help organizations orchestrate and govern data for AI workloads.

► **watsonx Orchestrate Agent Observability**

Added agent observability and governance capabilities to watsonx Orchestrate, providing real-time visibility into AI agents, lifecycle management, and automated policy enforcement for managing complex multi-agent systems with enterprise-grade guardrails.

Samsung

Samsung is positioned as a core provider in this year's DSML Platform Technology Value Matrix. The Brightics Machine Learning platform delivers an integrated environment for developing, deploying, and managing AI solutions with a modular architecture enabling tailored analytics workflows at enterprise scale. Samsung's Brightics Data Hub provides connectivity to diverse data sources with transformation tools for structured and unstructured data including text, images, time series, and tabular formats. A dual-track development approach combines Brightics AutoML for accelerated model creation with specialized environments for practitioners using Python, R, and deep learning frameworks including TensorFlow and PyTorch.

Samsung SDS's AX Center consolidates AI platform capabilities for managing the full AI agent lifecycle, reflecting the company's expansion from manufacturing AI into enterprise services.

The platform serves data engineers, data scientists, business users, and executives through differentiated interfaces, with model management capabilities tracking versions, parameters, and performance metrics. Access controls, audit trails, and drift detection support production governance. The FabriX platform and Brity Copilot extend Samsung's enterprise AI capabilities with large language model integration and agentic workflow execution. Samsung SDS has established a growing government and enterprise footprint in Korea, with the AX Center formalizing an AI agent lifecycle management practice across sectors.

Brity Copilot evolved from a collaboration assistant to an agentic platform with real-time interpretation across 60+ languages, growing.

Recent product updates and enhancements include:

▶ **FabriX Platform and OpenAI Partnership**

Samsung SDS became Korea's first OpenAI reseller partner for ChatGPT Enterprise. FabriX was applied to a pan-government AI infrastructure project serving 720,000 civil servants, with contracts secured across financial services and travel sectors.

▶ **Brity Copilot with AI Agent Capabilities**

Upgraded from a collaboration assistant to an agentic platform with Personal Agent voice conversations, speaker-separation-based real-time interpretation across 60+ languages, and automated task execution.

▶ **AX Center for Enterprise AI Agent Lifecycle**

Established to consolidate AI platform capabilities and manage AI agent lifecycle operations, build, deploy, control, monitor, across public, financial, and manufacturing sectors.