

IDC MarketScape: Worldwide Machine Learning Operations Platforms 2022 Vendor Assessment

Kathy Lange
Schubmehl

Raghunandhan Kuppuswamy

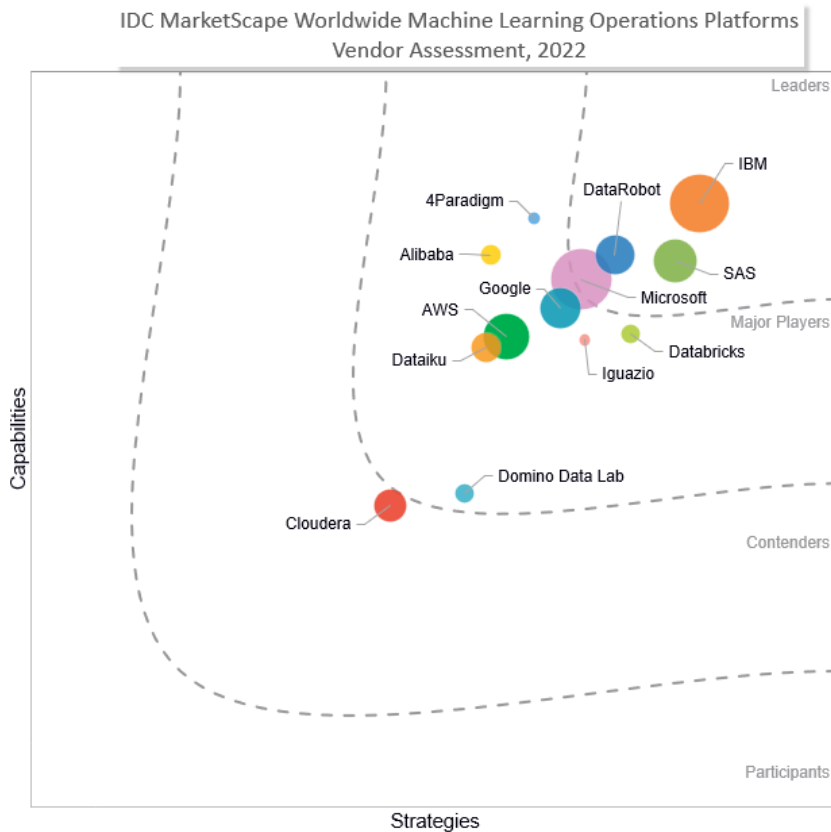
David

THIS IDC MARKETSCAPE EXCERPT FEATURES SAS

IDC MARKETSCAPE FIGURE

FIGURE 1

IDC MarketScape Worldwide Machine Learning Operations Platforms Vendor Assessment



Source: IDC, 2022

Please see the Appendix for detailed methodology, market definition, and scoring criteria.

IN THIS EXCERPT

The content for this excerpt was taken directly from IDC MarketScape: Worldwide Machine Learning Operations Platforms 2022 Vendor Assessment (Doc # US48325822). All or parts of the following sections are included in this excerpt: IDC Opinion, IDC MarketScape Vendor Inclusion Criteria, Essential Guidance, Vendor Summary Profile, Appendix and Learn More. Also included is Figure 1 and 2.

IDC OPINION

Overview

Enterprise organizations across industries are increasingly adopting artificial intelligence/machine learning (AI/ML) technologies to improve operational efficiencies, increase innovation, and improve user experience. According to IDC's Worldwide Semiannual Artificial Intelligence Tracker (1H22), the worldwide AI market, including software, hardware, and services, is forecast to grow 19.6% YoY in 2022 to \$432.8 billion and is expected to break the \$500 billion mark in 2023.

While the global adoption of AI/ML technologies is increasing, it is not without challenges. As more ML models are being moved to production, the end users often face challenges including cost, lack of automation, lack of expertise, and scale. These challenges impede an organization's ability to leverage ML capabilities to the fullest.

Challenges Implementing AI Solutions

In a recent IDC study (IDC's *AI Strategies BuyerView Global Survey*, CY21), respondents cited lack of automation, cost, and lack of expertise as top challenges during AI implementations.

Lack of Automation

While enterprises are leveraging AI capabilities to gain a competitive advantage by bringing innovations to market faster, they often are unable to accelerate as much as they would like to. Challenges such as excessive cost of model build and training, difficulties in moving models from experimentation to production, and complexity of managing multiple tools/platforms across stages of the ML pipeline impede AI acceleration. In the aforementioned IDC study, more than 55% of the respondents cited cost and lack of automation as their top challenges during AI implementations, slowing down their ability to bring out new capabilities faster.

Cost

About 55% of the respondents cited cost as one of the top challenges with AI implementations. A successful AI/ML implementation requires investments in acquiring the right expertise and infrastructure, data acquisition and management, and processes.

Lack of Expertise

Implementing AI capabilities requires expertise across various domains and technologies such as data engineering, data science, agile model development, and IT operations. Data engineering expertise is needed to prepare and manage data for building and training models. Data science expertise is required to build models, select the right learning algorithms to use, and fine-tune models to achieve necessary levels of accuracy. Since model training and model inferencing usually require accelerated hardware, allocating the right computational resources for these tasks during both experimentation and

production stages requires infrastructure operational expertise. Enterprise organizations may not have sufficient in-house expertise in these areas – so much that our IDC survey respondents cited lack of expertise as one of their top challenges with AI implementation.

What Is MLOps?

Machine learning operations (MLOps) technologies and processes help enterprise organizations overcome the aforementioned challenges. IDC's software taxonomy defines machine learning operations tools and technology as those that support model deployment, model management, and model monitoring. The key capabilities of an MLOps platform include and are not limited to those discussed in the sections that follow.

Model Serving

Model serving refers to the ability to serve model prediction capabilities through an endpoint/API. It is critical that the endpoints are secured (via HTTPS protocol) and can load balance the traffic.

Model Registry

The model registry offers a centralized store for machine learning models, a set of APIs, and user interface (UI) to collaboratively manage the full life cycle of the model. It enables tracking model lineage, model versioning, stage transitions, and annotations.

Model Tracking

Model tracking enables checking and tracking metrics such as model accuracy, precision, or a confusion matrix. It is usually part of the model validation phase and often involves logging parameters, code versions, metrics, and artifacts while executing model prediction and visualization.

Model Monitoring

Model monitoring involves checking models running in production for accuracy and detecting any performance drifts with production data. Model monitoring is not complete without the ability to notify of such drifts – so that appropriate action/troubleshooting can be triggered. Based on the issue, the model could be rebuilt and retrained, or the training data set may be modified. The most used notification mechanisms include SMS, email, and notifications on third-party communication tools such as Slack, ChatOps, or event triggers.

ML Pipeline and MLOps

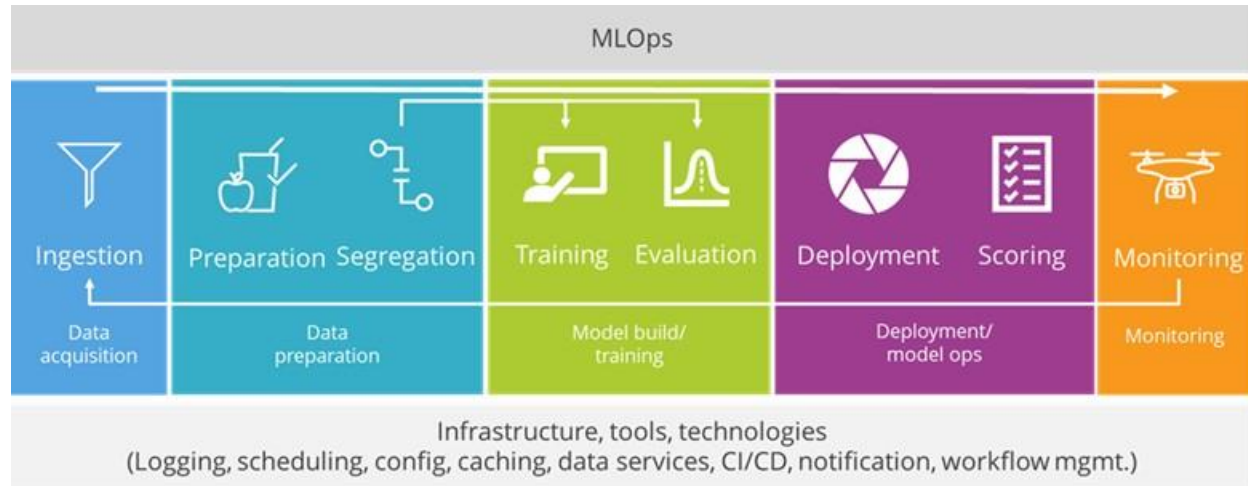
Informally, MLOps captures and expands on operational practices for software development to manage the unique challenges of machine learning and enables the practice of collaboration and communication between data scientists, data architects, business analysts, and operations professionals. As shown in Figure 2, it spans the entire ML life cycle, from experimentation to production, and is powered by newer sets of tools, libraries, and frameworks and aims to keep costs in check, simplifies management, and accelerates time to value.

A machine learning model goes through various stages across the machine learning pipeline, including data ingestion, data preparation, model exploration, model build/train, model scoring, application integration/model deployment, and monitoring. As these stages employ different personas, tools, and processes that most often operate in separate silos, they slow down the model velocity, with end users often finding it difficult to move models to production faster. Machine learning use cases also often involve multiple models, and enterprises tend to enable multiple use cases simultaneously.

MLOps stitches together different stages of the ML pipeline through automation. MLOps not only enables collaboration between data scientists and IT operators but also enables stronger collaboration between data scientists themselves for better model reuse. With MLOps tools and processes, IT operators can deploy, monitor, and troubleshoot models in production and trigger feedback loops back to the data scientists. Through automation and collaboration, MLOps also enables continuous delivery of machine learning models, thereby accelerating the pace of innovation.

FIGURE 2

ML Pipeline and MLOps



Note: For more information, see *Architecting Scalable Machine Learning Pipelines for AI-Enabled Enterprise Transformations* (IDC #US47992721, July 2021).

Source: IDC, 2022

IDC studies show that it takes about 290 days on average to fully deploy a model into production from start to finish. IDC also observes that as more models are getting deployed into production, end users are facing challenges with model performance, model drift, and bias. Such challenges and long development times slow down time to market, thereby impacting the organization's ability to bring out product innovations faster to the market.

Model velocity (MV) commonly refers to the time taken from the start to finish in the ML pipeline, from experimentation to production. Model velocity directly impacts an organization's ability to roll out new product features – the slower the model velocity is, the longer it takes to bring out the capability to market.

MLOps platforms can enable setting up scalable ML pipelines to manage, track, and monitor multiple models simultaneously. MLOps platforms can also help identify the issue and trigger workflows to model rebuild, retrain, or change data sets in case of model drift. Through these capabilities, end-to-end automation, and stronger collaboration between different personas, MLOps platforms can help accelerate model velocity.

The purpose of this document is to identify and evaluate vendors that offer MLOps platforms.

IDC MARKETSCOPE VENDOR INCLUSION CRITERIA

The inclusion criteria are as follows:

- The offering must be commercially available for use as a single product or a suite of services and for purchase by customers for at least one year (i.e., calendar year 2021).
- The offering must have the ability to deploy models, monitor model performance, detect model drift, enable building continuous integration/continuous delivery (CI/CD) pipelines of models, enable lineage tracking, enable model governance, and provide feedback loop into the ML pipeline. These capabilities can be available through APIs, UI, or both.
- The offering should include the following capabilities:
 - Model registry
 - Model tracking
 - Model deployment in on-premises, cloud, or edge locations
 - Monitor model performance
 - Model serving as endpoints
 - Detect and notify model drift
- The offering must have at least 10 commercial customers that used this product in calendar year 2021.
- The offering must be offered and available on a worldwide basis.
- The offering must have at least \$10 million in software/services revenue in calendar year 2021.

ADVICE FOR TECHNOLOGY BUYERS

Treat Models as Source Code

As with agile software development processes, IDC recommends treating machine learning models as source code to enable improved collaboration among data scientists and application developers, increased model reuse, and better tracking of model lineage. IDC studies show that end users typically leverage multiple models for a use case, and ML enables multiple use cases simultaneously. Treating models as source code enables users to efficiently build, share, and reuse models, thereby speeding up model development time. Model lineage also helps pinpoint the erring model to fix during troubleshooting.

Plan for Scale

IDC studies show that when end users find success with early AI/ML implementations, they tend to increase their investments in AI/ML initiatives for more use cases. As with data gravity, model gravity also increases the likelihood of employing more machine learning models in an enterprise organization. Enabling multiple use cases, with each potentially employing multiple models, increases the complexity of managing models significantly. IDC recommends planning for such a scale where machine learning pipelines move multiple models from experimentation to production simultaneously.

Choose the Right Vendor for Your Needs

This IDC MarketScape highlights many vendors, all of which are successful at various aspects of machine learning operations. Organizations should work to identify their needs for building, deploying, monitoring, and governing their machine learning models on a holistic basis. During the IDC

MarketScape evaluation, it became clear that many organizations are currently building and deploying machine learning models in several different departments, ranging from IT and data science groups to manufacturing, production, marketing, and research and development to just name a few. The needs of these various groups differ substantially and MLOps is still relatively nascent in terms of best practices and formalized groups within organizations. Make sure that your organization understands its machine learning life cycle and uses tools and products that will benefit that life cycle.

VENDOR SUMMARY PROFILES

This section briefly explains IDC's key observations resulting in a vendor's position in the IDC MarketScape. While every vendor is evaluated against each of the criteria outlined in the Appendix, the description here provides a summary of each vendor's strengths and challenges.

SAS

SAS is positioned in the Leader's category in the 2022 IDC MarketScape for worldwide machine learning operations platforms.

SAS Institute is an AI and analytics software vendor. Its core MLOps product is SAS Viya Model Manager, which supports SAS Viya machine learning offering. The product offers model registry, model life-cycle management, model tracking, model deployment (into batch, real-time REST API, and container locations), model serving (batch and REST API), model performance monitoring, model drift detection, and model drift monitoring through proprietary software.

Model Manager offers additional capabilities such as score testing, publish validation, model comparison, and custom process automation through workflows. These capabilities apply to models developed on the Viya platform, using SAS' graphical model pipeline tool or via code, as well as models developed in Python and R. Model Manager integrates with other SAS tools to incorporate model risk management, decision and rule management, and model deployment into real-time streaming devices. SAS has a significant partnership with Microsoft and can integrate with Microsoft Azure.

SAS Model Manager offers governance and monitoring capabilities that allow organizations to practice responsible AI. Facilities such as compliance and validation reporting provide enterprises with the ability to establish and follow best practices using Model Manager.

Quick facts about SAS include:

- **Year founded:** 1976
- **Headquarters:** SAS is headquartered in Cary, North Carolina, the United States, and is privately owned. SAS is currently planning an IPO in 2024.
- **Total number of employees:** 14,000
- **Deployment options:** Public cloud/private cloud/on premises/in containers/edge
- **Pricing model:** Pricing is on a per user/seat consumption basis depending on the service required. SAS Viya is now available through Microsoft Azure Marketplace. Customers will be charged on a pay-as-you-go basis.
- **Related products/services:** In addition to its core machine learning and analytics software, SAS offers related services including professional services, benchmarking, automated operations, and so forth.

Strengths

- **Flexible language and model support:** SAS Model Manager offers support for a wide variety of languages and platforms. It integrates with Git (including GitLab and GitHub), Python, R, MLflow, and Azure Machine Learning, as well as with container registries on Azure, AWS, GCP, and private Docker registries. Customers appreciated this flexibility and capability.
- **Strong governance and production capabilities:** As a longtime leader in the advanced analytics and AI space, SAS offers very strong monitoring and governance capabilities. Organizations use SAS Model Manager extensively to support the production runs of their machine learning models and their analytics models in a single environment.

Challenges

- **Pricing and flexibility:** While SAS is a well-known and successful analytics and AI software vendor, its pricing is not always seen as cost effective compared with other products and vendors in the market. This also extends to its Model Manager product. There are several open source and/or low-cost options for machine learning operations, and some companies are migrating to those platforms away from higher-cost products like that of SAS'.
- **Interoperability with other ML platforms:** While SAS has a strong ongoing partnership with Microsoft for machine learning operations software, it doesn't have the integrations with other machine learning software vendors that organizations would like to see. Many organizations have multiple commercial and open source machine learning systems in place. The ability to work with those platforms in a single environment, as SAS does with open source today, will be seen as increasingly important.

Consider SAS When

Consider SAS for machine learning operations if you are either an existing SAS client or if you're a multinational corporation looking for adaptable multi-domain support that can be deployed either on premises or via cloud. SAS Model Manager provides a wide range of related services and products that can assist with production of machine learning models in the enterprise.

APPENDIX

Reading an IDC MarketScape Graph

For the purposes of this analysis, IDC divided potential key measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis reflects the vendor's current capabilities and menu of services and how well aligned the vendor is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here and now. Under this category, IDC analysts will look at how well a vendor is building/delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the x-axis, or strategies axis, indicates how well the vendor's future strategy aligns with what customers will require in three to five years. The strategies category focuses on high-level decisions and underlying assumptions about offerings, customer segments, and business and go-to-market plans for the next three to five years.

The size of the individual vendor markers in the IDC MarketScape represents the market share of each individual vendor within the specific market segment being assessed.

IDC MarketScape Methodology

IDC MarketScape criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and specific vendors. IDC analysts tailor the range of standard characteristics by which vendors are measured through structured discussions, surveys, and interviews with market leaders, participants, and end users. Market weightings are based on user interviews, buyer surveys, and the input of IDC experts in each market. IDC analysts base individual vendor scores, and ultimately vendor positions on the IDC MarketScape, on detailed surveys and interviews with the vendors, publicly available information, and end-user experiences in an effort to provide an accurate and consistent assessment of each vendor's characteristics, behavior, and capability.

Market Definition

Machine learning operations (MLOps) tools and technology support model deployment, model management, and model monitoring. These capabilities include but are not limited to model deployment across different locations, serving models as endpoints, tracking model lineage, setting model performance metrics, monitoring model performance, troubleshooting model drift, enabling model governance, and providing feedback loop into different stages of ML pipeline. They support one or more data ingestion/preparation platforms, model build/train/inferencing platforms, and notification services.

LEARN MORE

Related Research

- *Market Analysis Perspective: Worldwide AI Life-Cycle Software, 2022* (IDC #US48545822, August 2022)
- *Worldwide AI Life-Cycle Software Forecast, 2022-2026* (IDC #US49433022, July 2022)
- *IDC Market Glance: AI and ML Life-Cycle Software, 4Q21* (IDC #US48447821, December 2021)
- *AI StrategiesView 2021 Standard: Banner Tables* (IDC #US47635721, April 2021)
- *Enterprise AI: An Architectural Shift Emerges* (IDC #US47602221, April 2021)
- *MLOps: Your Business' New Competitive Advantage* (IDC #US46643620, July 2020)

Synopsis

This IDC study evaluates vendors that offer MLOps technologies and capabilities. As customers are moving more models from experimentation to production, they need scalable ways to collaborate, operate, and operationalize machine learning models. To leverage the tremendous opportunities that this provides, IDC recommends leveraging MLOps tools, technologies, and methodologies to improve collaboration between data scientists and operational engineers, automate and accelerate model velocity, and provide for monitoring and governance of machine learning models in production.

"Top challenges that customers face with implementing machine learning initiatives in production include lack of expertise, cost, and lack of automation," said Kathy Lange, research director, AI Software research at IDC. "MLOps software and processes enable customers to overcome these challenges by improving collaboration between data scientists, application developers, and operational engineers; automating end-to-end model life-cycle management; and increasing model velocity."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com

Copyright and Trademark Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights. IDC and IDC MarketScape are trademarks of International Data Group, Inc.

Copyright 2022 IDC. Reproduction is forbidden unless authorized. All rights reserved.

