Agentic and Generative AI Buyers Guide

Software Provider and Product Assessment

ISG Research

Agentic and Generative AI

As artificial intelligence (AI) continues to evolve, there is an increasing need for it to move beyond merely providing scores or recommendations to going further and taking actionable steps. While scores and recommendations offer valuable insights, they often require decision-makers to translate the data into effective strategies. As enterprises invest in generative AI (GenAI), they recognize this need. ISG Market Lens Research shows one-third of enterprises (32%) are using the technology to address business process workflow management, representing the third-biggest initiative to date based on investment. And it is second on the list of GenAI use cases that will deliver the most benefit over the next two years.

ISG Research defines GenAI as the ability generate new, seemingly real content, including text, documents, images and other types of media. GenAI is often used as the technology behind chatbots. It is also used to generate code and to summarize documents. The "generative" aspect of GenAI stems from the fact that it creates new material based on instructions or prompts from the users. GenAI uses large language models (LLMs) to generate new material by predicting the next element of the response, whether that element is the next word of a reply, the next line of code in a software program or the next pixel in an image or video.

Because GenAI can generate responses to user prompts, it gives the appearance—especially when used in a chatbot—that, like a customer service agent, it can take other types of actions. A related set of capabilities called agentic AI has emerged to address this need. ISG Research defines agentic AI as the ability to take autonomous actions, involving multiple processes or systems, based on an understanding of the environment and the goals that should be achieved. These two sets of capabilities are closely related, with software providers combining generative and agentic AI in their offerings.

> **GenAI has sparked universal interest in AI across all industries, sizes of enterprises and among consumers.**

GenAI has sparked universal interest in AI across all industries, sizes of enterprises and among consumers. GenAI is finally delivering on the promise of natural language processing (NLP), making it easier to interact with computer systems and other technology. However, it is not without its flaws. Like all predictive technologies, it's not 100% accurate. For a variety of reasons, such as a lack of information or inaccurate information used to create the LLMs, the responses may not be correct. When inaccurate material responses are generated, it is referred to as a hallucination. If the training material is biased or offensive, the responses may be biased or offensive.

In order to increase the accuracy of responses and reduce hallucinations, enterprises started fine-tuning LLMs and using retrieval augmented generation (RAG). Fine-tuning a model is the process of taking a pre-trained model and training it further on a specific data set or domain.

RAG is the process of augmenting the prompts with additional information that was not used in the training process, such as internal documents or data. RAG also depends on vector processing to determine similarity between text submitted in a prompt and documents stored in a knowledge base in order to retrieve the most contextually relevant information. In addition, the process of constructing prompts that generated the best responses evolved into a discipline known as prompt engineering.

At the same time enterprises were trying to increase the accuracy of responses, they were also starting to question the costs associated with GenAI usage. While larger and larger models provided more accurate results and a wider variety of use cases, they were also driving up the costs of responding to user requests. Increasing model sizes also drive up the use of expensive GPU processing. As a result, enterprises began to explore the use of targeted, smaller models. The proliferation of models and model types led to the need to create and manage multiple models with model gardens or model catalogs. Now it is common for GenAI platforms to support and manage a variety of models depending on use cases. However, through 2027, due to a lack of tooling and governance, controlling the costs for GenAI deployments will remain a concern for one-third of enterprises, limiting their deployments and ROI.

GenAI markets also took a page from the traditional AI and machine learning (ML) market. The notion of machine learning operations (MLOps) was well established when GenAI exploded onto the scene. MLOps applies discipline to the process of developing and deploying ML models into production in order to provide repeatability, monitoring and governance. These same notions have been applied to the development and deployment of LLMs and are referred to as LLM operations or LLMOps. However, the market is still evolving.

The most recent changes in the GenAI market, which is still evolving rapidly, are based around the concept of agentic AI. As noted above, enterprises recognize the need to apply GenAI to business processes and workflow management. LLMs, the foundation models used to generate text, images and videos, are not necessarily the right models to generate actions. Foundation models must be trained on a set of actions and outcomes to effectively generate the right set of actions to achieve the desired goal. In some cases, large action models (LAMs) are being used in agentic AI processes, and in other cases, LLMs are being extended to incorporate actions and outcomes in the training data set.

As a result, modeling tools and evaluation tools need to be extended to support agents and their associated actions. Agents also need to be able to perceive their environment and

perform reasoning in order to match actions to goals. They must be able to support a variety of software applications, including pre-built and custom ones. And, unlike GenAI which is driven by a prompt, agents need to be able to execute autonomously.

To effectively apply agentic and GenAI in their organizations, enterprises need tooling and processes to support each of the needs outlined above. They must be able to develop, fine-tune, deploy and monitor models. They must be able to execute prompts and augment the responses to prevent hallucinations and inaccuracies. They must be able to incorporate agentic and GenAI into their business processes in the form of chatbots, assistants or autonomous agents. They must also be able to optimize the costs of using these systems to match their budgets.

Fortunately, the widespread interest and demand for agentic and GenAI has driven software providers to invest heavily in meeting these needs. Providers are racing to gather market share before it is too late; as a result, it is a very competitive market. However, it is also an immature and rapidly changing market. Many of the capabilities needed are still under development or in various stages of pre-release. The providers with more of these capabilities released and supported will be better positioned to meet today's enterprise needs.

The ISG Buyers Guide™ for Agentic and Generative AI evaluates software providers and products in the following key areas: agentic AI, GenAI, preparation of data used in AI processes, support for optimizing model execution, developer tooling and LLM operations.

This research evaluates the following software providers that offer products that address key elements of agentic and GenAI as we define it: Alibaba Cloud, Altair, Anthropic, Automation Anywhere, AWS, C3 AI, Cohere, Databricks, Dataiku, DataRobot, Domino Data Lab, Google Cloud, H2O.ai, Hugging Face, IBM, Microsoft, NVIDIA, OpenAI, Oracle, Palantir, Quantexa, Salesforce, SAP, ServiceNow, Snowflake, Teradata, UiPath and Weights & Biases.

Buyers Guide Overview

For over two decades, ISG Research has conducted market research in a spectrum of areas across business applications, tools and technologies. We have designed the Buyers Guide to provide a balanced perspective of software providers and products that is rooted in an understanding of the business requirements in any enterprise. Utilization of our research methodology and decades of experience enables our Buyers Guide to be an effective method to assess and select software providers and products. The findings of this research undertaking contribute to our comprehensive approach to rating software providers in a manner that is based on the assessments completed by an enterprise.

> **ISG Research has designed the Buyers Guide to provide a balanced perspective of software providers and products that is rooted in an understanding of business requirements in any enterprise.**

The ISG Buyers Guide™ for Agentic and Generative AI is the distillation of over a year of market and product research efforts. It is an assessment of how well software providers' offerings address enterprises' requirements for agentic and GenAI software. The index is structured to support a request for information (RFI) that could be used in the request for proposal (RFP) process by incorporating all criteria needed to evaluate, select, utilize and maintain relationships with software providers. An effective product and customer experience with a provider can ensure the best long-term relationship and value achieved from a resource and financial investment.

In this Buyers Guide, ISG Research evaluates the software in seven key categories that are weighted to reflect buyers' needs based on our expertise and research. Five are product-experience related: Adaptability, Capability, Manageability, Reliability, and Usability. In addition, we consider two customer-experience categories: Validation, and Total Cost of Ownership/Return on Investment (TCO/ROI). To assess functionality, one of the components of Capability, we applied the ISG Research Value Index methodology and blueprint, which links the personas and processes for agentic and GenAI to an enterprise's requirements.

The structure of the research reflects our understanding that the effective evaluation of software providers and products involves far more than just examining product features, potential revenue or customers generated from a provider's marketing and sales efforts. We believe it is important to take a comprehensive, research-based approach, since making the wrong choice of agentic and GenAI technology can raise the total cost of ownership, lower the return on investment and hamper an enterprise's ability to reach its full performance potential. In addition, this approach can reduce the project's development and deployment

time and eliminate the risk of relying on a short list of software providers that does not represent a best fit for your enterprise.

ISG Research believes that an objective review of software providers and products is a critical business strategy for the adoption and implementation of agentic and GenAI software and applications. An enterprise's review should include a thorough analysis of both what is possible and what is relevant. We urge enterprises to do a thorough job of evaluating agentic and GenAI systems and tools and offer this Buyers Guide as both the results of our in-depth analysis of these providers and as an evaluation methodology.

How To Use This Buyers Guide

Evaluating Software Providers: The Process

We recommend using the Buyers Guide to assess and evaluate new or existing software providers for your enterprise. The market research can be used as an evaluation framework to establish a formal request for information from providers on products and customer experience and will shorten the cycle time when creating an RFI. The steps listed below provide a process that can facilitate best possible outcomes.

1. Define the business case and goals.
   Define the mission and business case for investment and the expected outcomes from your organizational and technological efforts.
2. Specify the business needs.
   Defining the business requirements helps identify what specific capabilities are required with respect to people, processes, information and technology.
3. Assess the required roles and responsibilities.
   Identify the individuals required for success at every level of the enterprise from executives to frontline workers and determine the needs of each.
4. Outline the project's critical path.
   What needs to be done, in what order and who will do it? This outline should make clear the prior dependencies at each step of the project plan.
5. Ascertain the technology approach.
   Determine the business and technology approach that most closely aligns to your enterprise's requirements.
6. Establish software provider evaluation criteria.
   Utilize the product experience: Adaptability, Capability, Manageability, Reliability and Usability, and the customer experience in TCO/ROI and Validation.
7. Evaluate and select the technology properly.
   Weight the categories in the technology evaluation criteria to reflect your enterprise's priorities to determine the short list of software providers and products.
8. Establish the business initiative team to start the project.
   Identify who will lead the project and the members of the team needed to plan and execute it with timelines, priorities and resources.

The Findings

All of the products we evaluated are feature-rich, but not all the capabilities offered by a software provider are equally valuable to types of workers or support everything needed to manage products on a continuous basis. Moreover, the existence of too many capabilities may be a negative factor for an enterprise if it introduces unnecessary complexity. Nonetheless, you may decide that a larger number of features in the product is a plus, especially if some of them match your enterprise's established practices or support an initiative that is driving the purchase of new software.

Factors beyond features and functions or software provider assessments may become a deciding factor. For example, an enterprise may face budget constraints such that the TCO evaluation can tip the balance to one provider or another. This is where the Value Index methodology and the appropriate category weighting can be applied to determine the best fit of software providers and products to your specific needs.

Overall Scoring of Software Providers Across Categories

The research finds Google Cloud atop the list, followed by Oracle and IBM. Companies that place in the top three of a category earn the designation of Leader. Oracle has done so in six categories; Databricks and Google Cloud in four; Microsoft in three; AWS, Hugging Face, IBM and Teradata in one category.

The overall representation of the research below places the rating of the Product Experience and Customer Experience on the *x* and *y* axes, respectively, to provide a visual representation and classification of the software providers. Those providers whose Product Experience have a higher weighted performance to the axis in aggregate of the five product categories place farther to the right, while the performance and weighting for the two Customer Experience categories determines placement on the vertical axis. In short, software providers that place closer to the upper-right on this chart performed better than those closer to the lower-left.

Agentic and Generative AI

Overall

Providers	Grade	Performance
Google Cloud	B++	*Leader* 78.5%
Oracle	B++	*Leader* 77.4%
IBM	B++	*Leader* 76.0%
AWS	B++	75.1%
Databricks	B+	72.2%
Microsoft	B+	72.2%
UiPath	B	68.4%
ServiceNow	B	68.4%
Teradata	B	67.4%
Salesforce	B	66.3%
Dataiku	B	66.2%
Alibaba Cloud	B	65.8%
Snowflake	B	65.7%
DataRobot	B	65.0%
Automation Anywhere	B	64.8%
SAP	B	64.3%
Domino Data Lab	B	63.3%
Palantir	B	63.3%
NVIDIA	B	63.0%
Hugging Face	B-	62.2%
H2O.ai	B-	61.1%
OpenAI	B-	60.5%
Cohere	B-	56.3%
Weights & Biases	C++	56.1%
Altair	C++	55.7%
Anthropic	C++	50.9%
C3 AI	C+	49.3%
Quantexa	C	43.8%

ISG Source: ISG Research
Agentic and Generative AI Buyers Guide

The research places software providers into one of four overall categories: Assurance, Exemplary, Merit or Innovative. This representation classifies providers' overall weighted performance.

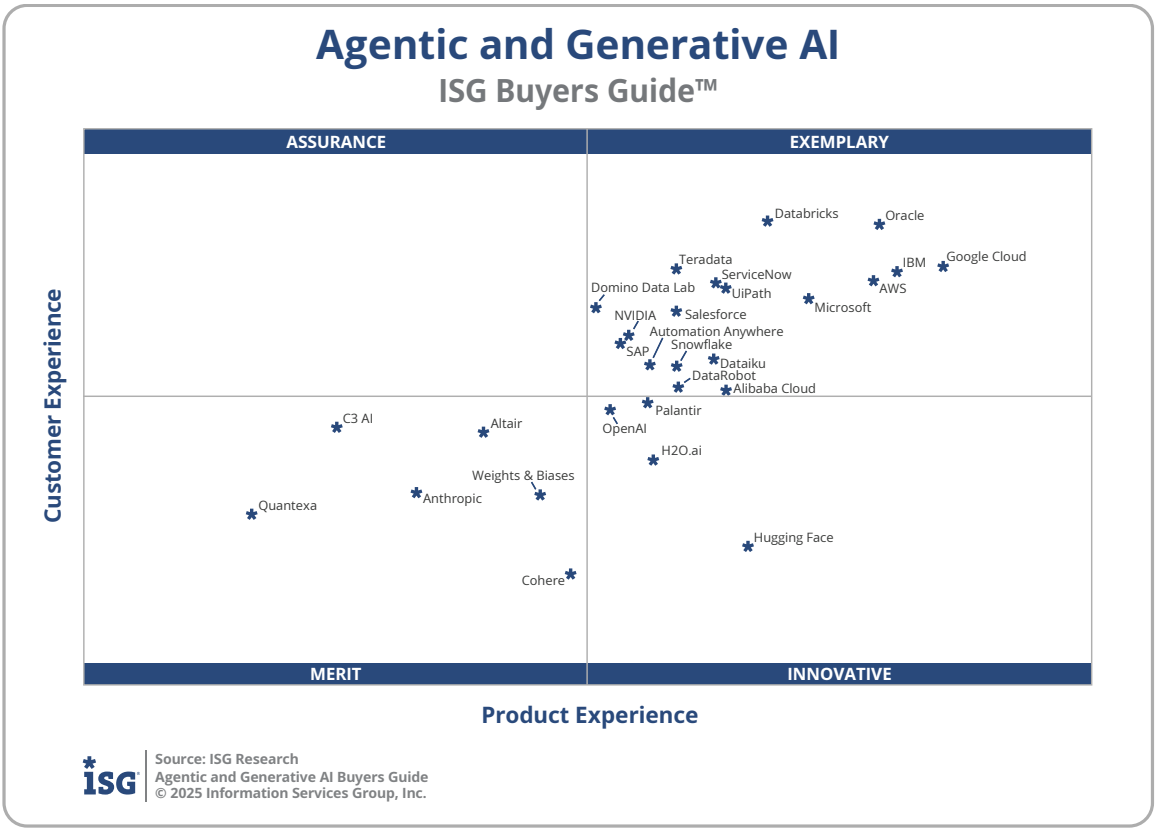

Exemplary: The categorization and placement of software providers in Exemplary (upper right) represent those that performed the best in meeting the overall Product and Customer Experience requirements. The providers rated Exemplary are: Alibaba Cloud, Automation Anywhere, AWS, Databricks, Dataiku, DataRobot, Domino Data Lab, Google Cloud, IBM, Microsoft, NVIDIA, Oracle, Salesforce, SAP, ServiceNow, Snowflake, Teradata and UiPath.

Innovative: The categorization and placement of software providers in Innovative (lower right) represent those that performed the best in meeting the overall Product Experience requirements but did not achieve the highest levels of requirements in Customer Experience. The providers rated Innovative are: H2O.ai, Hugging Face, OpenAI and Palantir.

Merit: The categorization of software providers in Merit (lower left) represents those that did not exceed the median of performance in Customer or Product Experience or surpass the threshold for the other three categories. The providers rated Merit are: Altair, Anthropic, C3 AI, Cohere, Quantexa and Weights & Biases.

We warn that close provider placement proximity should not be taken to imply that the packages evaluated are functionally identical or equally well suited for use by every enterprise or for a specific process. Although there is a high degree of commonality in how enterprises handle agentic and GenAI, there are many idiosyncrasies and differences in how they do these functions that can make one software provider's offering a better fit than another's for a particular enterprise's needs.

We advise enterprises to assess and evaluate software providers based on organizational requirements and use this research as a supplement to internal evaluation of a provider and products.

Product Experience

The process of researching products to address an enterprise's needs should be comprehensive. Our Value Index methodology examines Product Experience and how it aligns with an enterprise's life cycle of onboarding, configuration, operations, usage and maintenance. Too often, software providers are not evaluated for the entirety of the product; instead, they are evaluated on market execution and vision of the future, which are flawed since they do not represent an enterprise's requirements but how the provider operates. As more software providers orient to a complete product experience, evaluations will be more robust.

The research results in Product Experience are ranked at 80%, or four-fifths, of the overall rating using the specific underlying weighted category performance. Importance was placed on the categories as follows: Usability (10%), Capability (40%), Reliability (10%), Adaptability (10%) and Manageability (10%). This weighting impacted the resulting overall ratings in this research. Google Cloud, IBM and Oracle were designated Product Experience Leaders. While not a Leader, AWS was also found to meet a broad range of enterprise product experience requirements.

Agentic and Generative AI

Product Experience

Providers	Grade	Performance
Google Cloud	B++	*Leader* 60.4%
IBM	B+	*Leader* 58.0%
Oracle	B+	*Leader* 57.3%
AWS	B+	56.9%
Microsoft	B	53.8%
Databricks	B	51.5%
Hugging Face	B	50.4%
UiPath	B-	49.3%
Alibaba Cloud	B-	49.3%
ServiceNow	B-	49.2%
Dataiku	B-	48.9%
Snowflake	B-	47.9%
DataRobot	B-	47.9%
Salesforce	B-	47.9%
Teradata	B-	47.8%
H2O.ai	B-	47.0%
Automation Anywhere	B-	46.8%
Palantir	B-	46.6%
NVIDIA	B-	45.7%
SAP	B-	45.5%
OpenAI	B-	45.3%
Domino Data Lab	C++	44.7%
Cohere	C++	43.8%
Weights & Biases	C++	42.1%
Altair	C+	39.6%
Anthropic	C+	36.5%
C3 AI	C	33.6%
Quantexa	C-	28.8%

Source: ISG Research
Agentic and Generative AI Buyers Guide

Customer Experience

The importance of a customer relationship with a software provider is essential to the actual success of the products and technology. The advancement of the Customer Experience and the entire life cycle an enterprise has with its software provider is critical for ensuring satisfaction in working with that provider. Technology providers that have chief customer officers are more likely to have greater investments in the customer relationship and focus more on their success. These leaders also need to take responsibility for ensuring this commitment is made abundantly clear on the website and in the buying process and customer journey.

The research results in Customer Experience are ranked at 20%, or one-fifth, using the specific underlying weighted category performance as it relates to the framework of commitment and value to the software provider-customer relationship. The two evaluation categories are Validation (10%) and TCO/ROI (10%), which are weighted to represent their importance to the overall research.

The software providers that evaluated the highest overall in the aggregated and weighted Customer Experience categories are Databricks, Oracle and Google Cloud. These category leaders best communicate commitment and dedication to customer needs. While not Leaders, IBM and Teradata was also found to meet a broad range of enterprise customer experience requirements.

Agentic and Generative AI
Customer Experience

Providers	Grade	Performance
Databricks	A	*Leader* 17.7%
Oracle	A	*Leader* 17.6%
Google Cloud	A-	*Leader* 16.6%
IBM	A-	16.5%
Teradata	A-	16.4%
AWS	A-	16.3%
ServiceNow	B++	16.2%
UiPath	B++	16.1%
Domino Data Lab	B++	15.9%
Microsoft	B++	15.8%
Salesforce	B++	15.7%
NVIDIA	B++	15.5%
SAP	B++	15.4%
Dataiku	B++	15.2%
Automation Anywhere	B++	15.1%
Snowflake	B++	15.0%
DataRobot	B+	14.6%
Alibaba Cloud	B+	14.5%
Palantir	B+	14.4%
OpenAI	B+	14.2%
C3 AI	B+	13.8%
Altair	B	13.7%
H2O.ai	B	13.1%
Anthropic	B	12.7%
Weights & Biases	B	12.6%
Quantexa	B-	12.5%
Hugging Face	B-	11.6%
Cohere	C++	11.2%

ISG
Source: ISG Research
Agentic and Generative AI Buyers Guide

Software providers that did not perform well in this category were unable to provide sufficient customer case studies to demonstrate success or articulate their commitment to customer experience and an enterprise's journey. The selection of a software provider means a continuous investment by the enterprise, so a holistic evaluation must include examination of how they support their customer experience.

Appendix: Software Provider Inclusion

For inclusion in the ISG Buyers Guide™ for Agentic and Generative AI in 2025, a software provider must be in good standing financially and ethically, have at least $25 million in annual or projected revenue verified using independent sources, sell products and provide support on at least two continents and have at least 50 customers. The principal source of the relevant business unit's revenue must be software-related, and there must have been at least one major software release in the past 12 months.

To be included in the Agentic and Generative Buyers Guide, the software provider must enable the preparation of data used in agentic and generative AI processes and support optimizing model execution, generative AI, agentic AI, developer tooling and LLM operations.

The research is designed to be independent of the specifics of software provider packaging and pricing. To represent the real-world environment in which businesses operate, we include providers that offer suites or packages of products that may include relevant individual modules or applications. If a software provider is actively marketing, selling and developing a product for the general market and it is reflected on the provider's website that the product is within the scope of the research, that provider is automatically evaluated for inclusion.

All software providers that offer relevant agentic and GenAI products and meet the inclusion requirements were invited to participate in the evaluation process at no cost to them.

Software providers that meet our inclusion criteria but did not completely participate in our Buyers Guide were assessed solely on publicly available information. As this could have a significant impact on classification and ratings, we recommend additional scrutiny when evaluating those providers.

Products Evaluated

Provider	Product Names	Version	Release Month/Year
Alibaba Cloud	Platform for AI Model Studio	2025-03-28 N/A	March 2025 March 2025
Altair	AI Hub AI Studio	2025.0 2025.0	February 2025 December 2024
Anthropic	Claude	3.7	April 2025
Automation Anywhere	AI Agent Studio	36	March 2025
AWS	Amazon SageMaker AI	N/A	March 2025
C3 AI	C3 Agentic AI Platform C3 Generative AI	N/A N/A	April 2025 April 2025
Cohere	Cohere	API v2	April 2025
Databricks	Mosaic AI	N/A	April 2025
Dataiku	Dataiku	13.5.0	April 2025
DataRobot	Enterprise AI Suite	11.0.1	April 2025
Domino Data Lab	Domino Domino Cloud	6.0.4 N/A	April 2025 April 2025
Google Cloud	Vertex AI Platform	1	April 2025
H2O.ai	H2O AI Cloud	25.04.0	April 2025
Hugging Face	Enterprise Hub	0.31.0	April 2025
IBM	watsonx.ai	2.1.2	March 2025
Microsoft	Azure AI Foundry	N/A	April 2025
NVIDIA	NVIDIA AI Workbench	2025.04.01	April 2025
OpenAI	OpenAI	N/A	April 2025

Oracle	Oracle Cloud Infrastructure Generative AI Oracle Cloud Infrastructure Generative AI Agents Oracle Cloud Infrastructure Data Science	N/A N/A N/A	April 2025 April 2025 April 2025
Palantir	AIP	N/A	April 2025
Quantexa	Decision Intelligence Platform	2.7	April 2025
Salesforce	Data Cloud Agentforce	Spring '25 Spring '25	February 2025 February 2025
SAP	SAP Business AI (Joule) SAP AI Core	1.0 N/A	February 2025 April 2025
ServiceNow	Now Assist	Yokohama	April 2025
Snowflake	The Snowflake Platform	9.10	April 2025
Teradata	Teradata VantageCloud	N/A	April 2025
UiPath	UiPath Platform	2024.10	April 2025
Weights & Biases	W&B Weave W&B Models	0.68 0.68	April 2025 April 2025

Providers of Promise

We did not include software providers that, as a result of our research and analysis, did not satisfy the criteria for inclusion in this Buyers Guide. These are listed below as "Providers of Promise."

Provider	Product	Revenue	Capability	International	Customers
Clarifai	Clarifai	No	Yes	Yes	Yes
Crowdworks	Crowdworks AI	Yes	Yes	No	Yes
DeepSeek	DeepSeek Platform	No	Yes	No	Yes
EdgeVerve	AI Next	Yes	No	Yes	Yes
KNIME	KNIME Analytics Platform	No	Yes	Yes	Yes

About ISG Software Research and Advisory

ISG Software Research and Advisory provides market research and coverage of the technology industry, informing enterprises, software and service providers, and investment firms. The ISG Buyers Guides provide insight on software categories and providers that can be used in the RFI/RFP process to assess, evaluate and select software providers.

About ISG Research

ISG Research provides subscription research, advisory, consulting and executive event services focused on market trends and disruptive technologies. ISG Research delivers guidance that helps businesses accelerate growth and create more value. For further information about ISG Research subscriptions, please visit research.isg-one.com.

About ISG

ISG (Nasdaq: III) is a global AI-centered technology research and advisory firm. A trusted partner to more than 900 clients, including 75 of the world's top 100 enterprises, ISG is a long-time leader in technology and business services sourcing that is now at the forefront of leveraging AI to help organizations achieve operational excellence and faster growth. The firm, founded in 2006, is known for its proprietary market data, in-depth knowledge of provider ecosystems, and the expertise of its 1,600 professionals worldwide working together to help clients maximize the value of their technology investments.