



Agentic Support Services: Benchmarking Enterprise Software Customer Help

Key Findings

- Salesforce's Agentforce on Help is the #1 agentic help solution that incorporates enterprise-grade governance and advanced reasoning across all interactions.
- Salesforce also scored well in several other assessment categories, including accuracy/grounding, manageability/operations, and time to value.
- Most major SaaS vendors are highlighting agentic AI tools and offerings as a core feature and capability of their platform, with promises of greater efficiency, productivity, and effectiveness, though there is a large variation in the type and depth of guardrails and tools used to ensure the safety and accuracy of outputs.



Table of Contents

01/ Key Findings

03/ Background

04/ Proving ROI Via a Customer-Zero Approach

05/ Defining Agentic AI

06/ Competitive Set

06/ Adobe

06/ Atlassian

06/ AWS

06/ Google Cloud

06/ IBM

06/ Intercom

06/ Microsoft

06/ Oracle

06/ Salesforce

06/ SAP

06/ ServiceNow

07/ Research Background and Methodology

08/ Assessment and Results

09/ Results

10/ About Us

10/ About the Authors

10/ About The Futurum Group




Background

Bar none, artificial intelligence (AI) agents have quickly become the most talked-about technology in the enterprise technology market over the past 6-12 months. Regardless of the vendor, region, or industry, it is impossible to have any conversation that does not at least touch on AI agents, agentic tools, or agentic workflows.

However, like any new and exciting technology, there is a significant amount of hype, conjecture, and uncertainty about the capabilities of agentic AI. For enterprises that are considering the use of this technology, there is a clear delta between what capabilities are actually being delivered now versus what is on the horizon. Assessing the actual performance of AI agents is essential to cutting through the marketing hype and understanding how the technology can actually deliver real value.


A recent Futurum research study with 895 technology purchase decision makers and influencers indicates that despite the excitement around agentic AI, it is still not the leading purchase decision driver.

Figure 1: Top Criteria for Vendor Selection



Source: IH 2025 Enterprise Applications Decision Maker Survey, Futurum Research, June 2025

Figure 2: Current Approach to Agentic AI



Q: Which statement best describes your organization's current approach to agentic AI (AI systems that can take autonomous actions and make decisions on behalf of users)? (n = 838)

Source: Futurum Research, June 2025

That said, among the largest enterprises (those with >\$25 billion in annual revenue), 15% of survey respondents cite agentic AI as the top criteria for vendor selection (along with features, also cited by 15%). For other company sizes, agentic AI is only cited as a key criterion in the 7%–10% range. Large enterprises seek AI that not only generates responses, but takes actions autonomously. Vendors should target this segment with intelligent automation, AI-based reasoning, and agentic orchestration capabilities.

Proving ROI via a Customer-Zero Approach

Agentic AI is being marketed and positioned as a catalyst for dramatically changing the way organizations operate, shifting dull and repetitive tasks away from human workers, while automating complex, data-driven tasks and workflows. In essence, the goal with AI agents is to drive more efficiency, productivity, accuracy, and speed across the entire organization, while reducing costs.

However, these promises can often seem too good to be true. To that end, many vendors are deploying agentic technology within their own help or support offerings as a so-called “Customer Zero.” This approach can help a vendor refine and improve the technology on a continuous basis, while also demonstrating the effectiveness of the technology to the industry. Generally speaking, there will be a sufficient volume and velocity of interactions upon which to deploy the technology, and the disparity of inquiry or interaction types can further illustrate the flexibility and scalability of agentic technology.

Most importantly, by serving as “Customer Zero,” vendors can develop best practices they can share with prospects and customers as they begin to adopt AI agents and tools. Assessing vendors’ AI agents and tools deployed on their own help platforms for the benefit of their customers can provide key insight into the vendor’s progress and achievements with agentic technology. But first, it’s important to clearly define agentic AI to ensure a fair and level-playing field for comparison.



Defining Agentic AI

Futurum uses the following definition of an “agentic AI interaction” when it assesses whether the technology is being utilized by a vendor:


An “agentic AI interaction” is an instance where an AI system can operate without constant or persistent human guidance, can adapt its approach based on changing conditions or variables, and can initiate actions or make decisions based on its understanding of the situation. Unlike pre-programmed bots, agentic AI is capable of learning from experiences, refining its behavior over time, and optimizing itself over time to improve its performance.

Essentially, the definition focuses on four components: ability to operate independently, ability to adapt to changing conditions, can initiate actions or make decisions on its own, and can learn from experiences and self-optimize.

Most major SaaS vendors are highlighting agentic AI tools and offerings as a core feature and capability of their platform, with promises of greater efficiency, productivity, and effectiveness. Under the hood, most SaaS-based AI agents are large language model (LLM)-driven workflows with orchestration. A typical architecture uses a large language model (like GPT) as the cognitive engine, augmented by tools and data connectors.

For example, an agent might use retrieval-augmented generation (RAG) to fetch relevant documents, then let the LLM write a response, and finally call an API to enact an action (such as creating a support ticket or case, sending an email, or routing a customer to a human agent). These systems often incorporate planning algorithms or frameworks to break tasks into steps, but they do not continuously update their knowledge base or model weights on the backend.

However, at this point, human oversight remains critical to ensuring that today’s agents work as designed, operating safely and accurately. Total agentic autonomy is risky in business settings. Instead, vendors have leaned on explainability and supervision of agents, and as such, the fourth criterion used to determine an agentic system – self-learning or optimization – is yet to be enacted by vendors, and is not being considered as part of this agentic assessment criteria. As the technology is tested and improved, self-optimizing agents likely will become table stakes in certain use cases where speed and agility are core requirements.



Competitive Set

The following vendors' agentic AI offerings were evaluated based on the criteria and methodology laid out in the next section.

	Support Offering Link: helpx.adobe.com/support.html Agent Stack: Adobe Sensei + Adobe Firefly + Custom LLM Assistants (emerging) Notes: Adobe's Help Center offers extensive support for its apps, services, accounts, plans, and enterprise/team users, leveraging AI for issue resolution.
	Support Offering Link: https://support.atlassian.com/ Agent Stack: Atlassian Intelligence Notes: Embedded AI agents in Jira/Confluence answer product and technical questions; focused on developer/IT productivity support.
	Support Offering Link: console.aws.amazon.com/support/home Agent Stack: Amazon Q (Business + Developer) Notes: Q has been embedded into the console and support tiers; used for technical questions, account issues, and cloud service guidance.
	Support Offering Link: cloud.google.com/support Agent Stack: Google Agentspace + Vertex AI Agents Notes: Supports developers and admins with MCP-connected agents capable of context-aware multi-turn assistance and action.
	Support Offering Link: www.ibm.com/mysupport Agent Stack: watsonx Assistant + Discovery Notes: AI agents assist with Tier 0/I issue resolution and knowledge guidance; focused on enterprise tools and services.
	Support Offering Link: www.intercom.com/help Agent Stack: Fin AI Agent Notes: Fully autonomous agent resolving common user and account questions; smaller volume, but high agent autonomy.
	Support Offering Links: support.microsoft.com / azure.microsoft.com/en-us/support/ Agent Stack: Dynamics 365 Copilot + AI Service Agents Notes: Copilot agents integrated across Office, Azure, Dynamics; handles high-volume end-user and developer queries using action-capable agents.
	Support Offering Link: support.oracle.com Agent Stack: Oracle Digital Assistant + Fusion AI Notes: AI agents assist with navigating KBs and support requests; used in support for Fusion Apps, OCI, and NetSuite.
	Support Offering Link: help.salesforce.com Agent Stack: Agentforce Notes: Full agent orchestration with prebuilt actions, stateful agents, and escalation routing; used for CRM, platform, and Trailhead support.
	Support Offering Link: support.sap.com and help.sap.com Agent Stack: SAP Joule AI Copilot is actively used in SAP for Me Notes: Proactive case status updates, knowledge suggestions, multi-step workflows for common requests.
	Support Offering Link: support.servicenow.com Agent Stack: Now Assist + Virtual Agent Notes: AI agents used more heavily internally and for ITSM automation; public customer-facing agentic use is modest but growing.


Research Background and Methodology

To assess each vendor, Futurum conducted a multi-step evaluation process of each vendor's agentic help offering, using a combination of publicly available documentation (10-Ks, earnings call transcripts, marketing materials, website scans, industry articles, etc.), third-party web traffic sites, and, where available, direct information requests with vendors.

The information provided by vendors was cross-checked against public information, and harmonized based on a team of Futurum analysts' own assessments of each solution. The data was input into Futurum's model, which then created a data-based ranking incorporating all of the inputs.

Table 1. Agentic AI Vendor Assessment Criteria

Attribute or Metric	Description
Gross Monthly Business Help Traffic:	Overall business help traffic to the help site
Estimated Monthly Help Transactions:	Overall monthly help inquiries/requests
Estimated Percent Agentic Transactions:	Percentage of agentic transactions/total transactions
Estimated Monthly Agentic Transactions:	Numerical count of agentic transactions
Governance Layer:	Presence of a governance layer
Governance Layer Percent Coverage of Agentic AI:	Percentage of agentic transactions that utilize a governance layer
Reasoning Layer:	Presence of a reasoning or thinking layer
Reasoning Engine Percent Coverage of Agentic AI:	Percentage of agentic transactions that utilize a reasoning layer
Estimated Monthly Full Agentic Transactions:	Total monthly count of full agentic transactions
Conversation Volume:	Futurum Analyst Assessment/Ranking (1–10)
Successful Resolutions:	Futurum Analyst Assessment/Ranking (1–10)
Automation Depth:	Futurum Analyst Assessment/Ranking (1–10)
Task/Action Handling:	Futurum Analyst Assessment/Ranking (1–10)
Integration Ability (# of Integrations):	Futurum Analyst Assessment/Ranking (1–10)



Assessment

When assessing the competitive agentic solutions, Futurum developed two models to incorporate the data collected from public sources, vendors, and its own information gathering (experience, vendor briefings, etc.). At least one business use case was known to be implemented using AI agents by each vendor, though not necessarily for 100% of such requests.

The first model, which is an overall ranking model, incorporates the full business web traffic (subtracting any consumer help traffic) of each service, and reduces it to the number of help transactions (2.5 visits/searches before a problem is resolved on average) and then further assess the number of agentic transactions served up. It does NOT account for if a safety/trust layer is used or a reasoning engine.


Given that accuracy and safety is a prime component for enterprise use of agentic AI, Futurum then developed a Full Agentic Ranking Model. This model incorporates not only all of the same data included in the first model, but also credit for the use of a safety/governance layer and the use of a reasoning engine. A safety/governance layer ensures that outputs are relevant, free from toxicity, and acceptable for use in a business context. Meanwhile, a reasoning engine is designed to help an agent understand the context and semantic meaning of a help request, which, in turn, improves the accuracy of the agentic response. Agentic systems with low accuracy are generally more expensive to operate and frustrate users, and often wind up adding volume to traditional support channels.

Results

Agentic AI is still in its infancy; most solutions in the market today technically do not meet Futurum's (admittedly) stringent requirements, particularly around the self-optimizing functionality of a true AI agent. That said, two key components of agentic technology designed for use within or by enterprises are robust agentic governance and the use of a reasoning engine. These elements are critical to the definition of a true agentic system, as they ensure the technology is safe, robust, and accurate enough for wide-ranging business use cases.

As such, based on Futurum's assessment of agentic help systems, Salesforce's Agentforce on Help is the leading enterprise-grade agentic help platform that incorporates a governance layer and a reasoning engine, handling the largest monthly volume of agentic interactions, recently surpassing the milestone of handling 1 million support requests.

Figure 3. Technology Company Customer Help Platforms By Agentic Volume



While other vendors may have a greater number of agentic interactions, those generally use a more basic model of an agentic tech stack, in which requests are simply passed through an LLM that is grounded to a corpus of data.

Salesforce also scored well in other assessment categories, including accuracy/grounding, manageability/operations, and time to value. Accuracy and grounding are key elements of ensuring that an agentic system can avoid hallucinations. Agentic help solutions that do not ensure that agentic AI systems ground themselves in vetted data when making decisions or taking actions are unlikely to be trusted and used.

Agentic management is integral to ensuring that agents are working as they have been designed, and will become even more critical as these proliferate and interact with one another to complete more complex help-related tasks and actions that may include dependencies.

Finally, time to value is a measure of how quickly customers benefit after adoption of the technology. AI agents that are too complex or confusing to use are often difficult to scale beyond pilots, and may not get the traction required to justify investment and support in a production environment.

Important Information About This Report

AUTHORS

Dion Hinchcliffe

Vice President & Practice Lead,
Digital Leadership & CIO | The Futurum Group

Keith Kirkpatrick

Research Director, Enterprise Software
& Digital Workflows | The Futurum Group

Nick Patience

Vice President & Practice Lead,
AI Platforms | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

This document, including any supporting materials, is owned by The Futurum Group. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of The Futurum Group.

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.



ABOUT SALESFORCE

Salesforce is a global leader in customer relationship management (CRM) software, enabling businesses to connect with their customers in innovative ways. Its platforms support a variety of business functions, including sales, marketing, customer service, and analytics, helping organizations manage customer interactions and data efficiently. Known for its emphasis on innovation, Salesforce integrates artificial intelligence tools, such as Einstein and Agentforce, to enhance decision-making and automation. Serving businesses of all sizes across industries, Salesforce has become a central player in digital transformation.

Futurum

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum

CONTACT INFORMATION: The Futurum Group LLC | futurumgroup.com | (833) 722-5337