



Market Brief

AWS Doubles Down on Agentic AI; Announces Amazon Bedrock AgentCore

Key Enterprise Insights from AWS Summit New York, 2025

July 2025

Author: Yashashree Vivarekar

Executive Summary

AWS has recently announced the launch of Amazon Bedrock AgentCore to help organisations deploy and operate AI agents at enterprise scale along with its core services. It is compatible with any model and framework, thereby making it suitable for enterprises to move agents from prototypes to applications that can scale to large number of end-users. Moreover, it is also focused on reducing the complexities associated with security, integration, and compliance.

CIOs should evaluate its fit within multi-cloud environments and existing tech stacks, along with considerations around cost/value, and security.

Note – The term Amazon Bedrock AgentCore is also used interchangeably with Amazon Bedrock AgentCore, Bedrock AgentCore and AgentCore in the below sections. It refers to the same platform, that was introduced by AWS at its AWS Summit New York 2025. The platform supports the deployment and operation of AI agents at an enterprise level. Thus, this interchangeable use is intended to maintain clarity while ensuring a concise and consistent narrative.

Introducing Amazon Bedrock AgentCore

Amazon Web Services (AWS) introduced its **Amazon Bedrock AgentCore at the AWS NY Summit 2025**. It is a comprehensive set of enterprise-grade services that enables developers to deploy and operate AI agents at scale using any framework and model, that is hosted on Amazon Bedrock or elsewhere.

Introduction of Modular Stack to Accelerate Enterprise AI Adoption

Amazon Bedrock AgentCore is developed as a modular, production-ready platform that brings together key services, that are needed to deploy advanced, high-performing AI agents. Some of the key components of Amazon Bedrock AgentCore include the following –

- **AgentCore Runtime** – It is a foundational layer that delivers secure, serverless execution with complete isolation. Additionally, it is compatible with several agent frameworks. Moreover, these are designed for the dynamic, long-running nature of AI agents.
- **AgentCore Memory** – It is one of the key components that handles short-term conversational context and also handles long-term knowledge retention. With this feature, the agents store and access memories with minimal code, thereby leading them to learn and adapt across sessions.
- **AgentCore Identity** – It is an enterprise-grade security that ensures smooth integration with identity providers including Microsoft Entra ID, Amazon Cognito, and Okta. Additionally, it also handles API keys and OAuth credentials with complete security, thereby providing the needed authorization across various systems.
- **AgentCore Interpreter** – It creates an environment where agents can execute JavaScript, TypeScript, and Python code. This enables complex data analysis and workflow automation while maintaining security through VM-level isolation.
- **AgentCore Gateway** – It is an integration layer that transforms existing APIs, lambda functions, and services into agent-compatible tools. It supports the Model Context Protocol (MCP) and provides context-aware semantic search for tool discovery.
- **AgentCore Browser** – It is a cloud-based browser that allows agents to interact with web applications. It offers sub-second latency and compute-based session isolation, enabling agents to perform complex web tasks securely.
- **AgentCore Observability** – It offers real-time monitoring and debugging capabilities. Additionally, it is also compatible with OpenTelemetry, thereby helping teams in the identifying and addressing the challenges in agent development, such as understanding what the agent is doing and the reason behind the same

Why this matters? Analyst View (1/2)

Amazon Bedrock AgentCore addresses several concerns related to security, performance, and integration, thereby helping companies to adopt and manage AI agents more efficiently at scale. It is imperative for technology decision makers to consider the following:

- **Platform fit: Integration, interoperability, and strategic alignment in a multi-cloud world.** The platform is focused on positioning itself as a primary option for building and managing AI agents at scale. One of the primary concerns for CIOs and CTOs is how Bedrock AgentCore (or any other platform) would integrate and align within their existing technology stack to support the broader AI strategy. Additionally, several companies operate with a combination of cloud providers, AI models, and internal applications, some of which already include agent capabilities. This raises several questions around platform standardization and the ease of integrating AI agents across diverse systems. AgentCore is positioned as a solution, that offers compatibility with several agent frameworks and tools. It supports open-source agents and emerging interoperability standards, making it easier to connect across different systems.
- **Scrutinizing security, governance, and reliability in AgentCore.** Security, governance, and reliability remain key factors when evaluating Bedrock AgentCore. Technology leaders are assessing whether the platform aligns with their compliance requirements and can support deployment at scale. According to AWS, BedRock AgentCore is built for enterprise use, offering secure management of user sessions and credentials, along with the ability to scale across a large user base.
- **Addressing ops needs from experimentation to production.** For the companies integrating and using AI agents, Bedrock AgentCore offers several features that are focused on addressing operational challenges. For instance, the infrastructure of Bedrock AgentCore supports longer-running tasks while maintaining session isolation and security. Moreover, the agents are designed to retain relevant information across interactions, such as user preferences or key data points owing to its built-in memory capabilities. It also integrates with identity providers and manages access tokens securely, enabling safe interaction with internal systems. These capabilities are designed to support the transition from experimentation to more stable, production-level deployments.

Why this matters? Analyst View (2/2)

- **Simplifying Agentic adoption.** AgentCore is also designed to simplify adoption for organizations that are in the early stages of exploring and implementing AI agents.
 - Additionally, a section in the AWS marketplace offers ready-made agents and tools that can be deployed with minimal setup.
 - The Gateway service enables companies to adapt their existing APIs for use with AI agents with minimal development effort, and supporting faster experimentation.
 - With the presence of additional built-in tools, including a web-browser and code execution environment, allows agents to handle tasks without the requirement of extensive custom setup.

Concluding Remarks: Considerations for Tech Buyers

AgentCore envisions to simplify and accelerate Agentic AI adoption for enterprises—but for technology buyers, real and high-impact value realization will depend on a host of factors such as platform alignment & integration, costs, security and compliance. CIOs must explore avenues to balance short-term quick wins with a carefully crafted and orchestrated Agentic architecture.

- Technology decision makers must ensure that the platform does not lead to a vendor lock-in and is interoperable with existing or any potential Agentic AI platforms that the organization may want to leverage in the future.
- Customization efforts around integration of the platform with existing systems must be carefully assessed to ensure appropriate value is extracted from application of such efforts.
- Ensure AgentCore aligns with your organization's security, privacy, and compliance requirements.
- Move beyond the buzz around Agentic AI and identify high-impact use cases that align with your organization's business priorities and/or challenges.

Check Our Past Reports



Top Agentic AI Service Providers PeMa Quadrant 2025



State of Data Engineering in India – 2025



GCC Watchlist 2025: Top Fortune 500 U.S. Companies Likely to Establish in India

Upcoming Reports

Agentic AI Industry Trends: H1 2025 Update

Data Science Service Providers PeMa Quadrant 2025

PeMa Quadrant 2025– GCC Service Providers in India

Analyst

Yashashree Vivarekar

Research Analyst– Analytics & AI

Yashashree is a Research Analyst at AIM Research, where she focuses on Artificial Intelligence—including Generative AI and Agentic AI—as well as the evolving landscape of AI startups and Global Capability Centers (GCCs). With over four years of experience in the market research industry, she brings a strong foundation in both secondary research and client-driven engagements, delivering actionable insights to technology providers, enterprises, and innovation hubs.

At AIM Research, Yashashree plays a key role in developing strategic research reports and advisory content, with a focus on AI trends, funding and acquisition activity, and GCC transformation initiatives. Her work supports stakeholders in navigating a dynamic AI ecosystem and identifying emerging opportunities in both the technology and services sectors.



yashashree@aimresearch.co



<https://www.linkedin.com/in/yashashree-vivarekar-a1b60a103/>



AIM India

1st Floor, Sakti Statesman, Marathahalli — Sarjapur Outer Ring Rd, Green Glen Layout, Bellandur, Bengaluru — 560103

AIM Americas

2955, 1603 Capitol Avenue, Suite 413A, Cheyenne, WY, Laramie, US, 82001

www.aimresearch.co

info@aimresearch.co