



MongoDB Launches Atlas Agent Engine to Put AI Agents in Production Without a New Stack

September 29, 2026

Atlas Agent Engine unifies governance, memory, and retrieval, ending the choice between a stitched-together stack and being locked into one vendor's model and cloud

NEW YORK, Sept. 29, 2026 /PRNewswire/ -- **MongoDB, Inc. (NASDAQ: MDB)** today launched Atlas Agent Engine, a unified execution, memory, and governance layer for production AI agents, at its Investor Day at the Nasdaq MarketSite in New York City. While AI agents prove their value quickly in proof-of-concept testing, bringing them into production remains a major bottleneck. Doing so requires accurate retrieval, persistent memory, and enterprise-grade security and governance. Without a single platform, engineering teams must stitch together disparate tools that break every time underlying models or frameworks evolve. That's what Atlas Agent Engine is designed to solve for.



"Investigating unusual activity in our payment network today means our analysts stitching together data from multiple systems by hand, often under time pressure. We're excited about the potential for an intelligent agent, built on MongoDB's Atlas Agent Engine, to shrink the time between a problem emerging and our team acting on it, giving our analysts more time to focus on the judgment calls that matter most," said Amar Akshat, SVP of Architecture, Paysafe.

Atlas Agent Engine is available today in public preview. New and existing Atlas customers can get started at agentengine.mongodb.com.

Atlas Agent Engine gives teams a modular way to put agents into production. Retrieval is powered by MongoDB Voyage AI, whose embedding and reranking models rank among the top performers on RTEB, a benchmark built to reflect real enterprise retrieval instead of academic datasets. Customers can adopt the memory and governance layers independently or with the runtime, using existing models and frameworks they know. Atlas Agent Engine is available today in public preview, with consumption-based pricing for Atlas Agent Runtime and Atlas Agent Memory. Usage draws on customers' existing Atlas commitments, so adoption extends infrastructure already in place rather than requiring a new contract.

"Context is the critical success factor in successfully using agents for application development. Enterprises are currently struggling to assess, integrate and manage information across multiple systems to enable an ontology for autonomous agentic work." said James Governor, co-founder of RedMonk. "MongoDB Atlas Agent Engine is designed to bake governance into agentic app development with a single platform for memory and identity."

Atlas Agent Engine is built with the ecosystem enterprises already trust. Frontier Labs bring the best available models directly to where enterprise data already lives. System integrator partners bring the industry expertise and delivery experience customers rely on for agent deployments.

"Atlas Agent Engine brings the enterprise-ready capabilities, context, and constraints needed to help AI agents deliver real-world impact. Combined with Accenture's governance, architecture, and deep industry expertise, it creates a powerful foundation for accelerating AI transformation and delivering outcomes at scale. Our shared commitment to delivery and customer success makes this partnership particularly strong," said Ram Ramalingam, Global Lead, SW Engineering & Head of RDE Accenture.

Built to get agents into production

Enterprises building agents often run into the same three challenges: actions nobody can govern, agents that forget, and lock-in to a single model or framework. Atlas Agent Engine solves all three, grounded in the same operational platform that more than 70,000 customers already run on, with enterprises like Paysafe already building toward production.

"Organizations that want to put agents in production are being forced into a false tradeoff: either adopt one vendor's runtime and accept being locked into a model and cloud, or piece together a framework and manage governance and memory on their own," said Pablo Stern-Plaza, Chief Product Officer, AI and Emerging Products, MongoDB. "With the launch of Atlas Agent Engine, that false tradeoff ends today. Enterprises get the real-time context their agents need, with governance and security built in from the start, and the freedom to run any model, any framework, and on any cloud. We didn't want to ask customers to predict the future. We wanted to build something that works no matter what they choose."

Governed by default. Most platforms handle identity, audit, guardrails, and cost controls as separate systems teams have to stitch together themselves. Atlas Agent Engine puts it all behind one control plane: every action is logged against a real identity, human or agent, and governed by policy that can't be quietly switched off. Governance is built in, not bolted on after launch. So when someone asks what an agent did and who authorized it, the answer takes seconds, not weeks. And because governance, memory, and retrieval run as one system instead of stitched-together services, there's less to secure and fewer places for things to break.

Memory and retrieval built in. Without built-in memory, agents start every conversation from zero, and teams end up rebuilding memory infrastructure for every new agent. Atlas Agent Engine builds memory into the platform itself, using Voyage AI embeddings and MongoDB's native retrieval, so

agents get more accurate while spending fewer tokens.

Open design. Standardizing on one model, cloud, or framework is one of the riskiest infrastructure bets a leader can make in a market that moves this fast. Atlas Agent Engine is neutral across AI models and frameworks. Because it's built on open standards like MCP and A2A, changing course later only takes a configuration change rather than an expensive rebuild. It will also run across any cloud, self-managed or even a laptop, so the same agent works everywhere without rebuilding cloud by cloud. Atlas Agent Engine adds governed execution, memory, and cost control on top of what teams already run, rather than asking them to replace it.

MongoDB is committed to giving customers the openness and security they need. We are joining the Linux Foundation's Open Secure AI Alliance and Agentic AI Foundation to drive open software and standards for secure, interoperable agents—so organizations can move agents into production with the flexibility and control to scale as their needs evolve.

Atlas Agent Engine, MongoDB 9.0, and Atlas Infinite, also announced today, build on each other: MongoDB 9.0 strengthens the foundation every MongoDB customer already runs on, Atlas Infinite removes the limits on how that foundation can scale, and Atlas Agent Engine puts AI agents to work on top of both, governed and grounded in real-time data. Together with Voyage AI's industry-leading embedding and retrieval models, already generally available, this extends MongoDB's intelligent data platform for the AI era. To learn more about all three new launches and everything else MongoDB announced at Investor Day, please visit <https://www.mongodb.com/company/newsroom>.

About MongoDB

MongoDB is the intelligent data platform for the AI era. Headquartered in New York, MongoDB helps organizations move faster, build smarter, and simplify complex architectures on a unified platform for operational data, search, real-time analytics, and AI-powered retrieval, coupled with MongoDB Voyage AI's industry-leading embedding and reranking models. The platform is available as a fully managed service with MongoDB Atlas or self-managed with MongoDB Enterprise Advanced or Community Edition. Millions of developers and more than 70,000 customers, including approximately 75% of the Fortune 100, rely on MongoDB for their most important customer-facing applications. To learn more, visit mongodb.com.

Forward-Looking Statements

This press release includes certain "forward-looking statements" within the meaning of Section 27A of the Securities Act of 1933, as amended, or the Securities Act, and Section 21E of the Securities Exchange Act of 1934, as amended, including new capabilities announced at MongoDB.local NYC Investor Day 2026. These forward-looking statements include, but are not limited to, plans, objectives, expectations and intentions and other statements contained in this press release that are not historical facts and statements identified by words such as "anticipate," "believe," "continue," "could," "estimate," "expect," "intend," "may," "plan," "project," "will," "would" or the negative or plural of these words or similar expressions or variations. These forward-looking statements reflect our current views about our plans, intentions, expectations, strategies and prospects, which are based on the information currently available to us and on assumptions we have made. Although we believe that our plans, intentions, expectations, strategies and prospects as reflected in or suggested by those forward-looking statements are reasonable, we can give no assurance that the plans, intentions, expectations or strategies will be attained or achieved. Furthermore, actual results may differ materially from those described in the forward-looking statements and are subject to a variety of assumptions, uncertainties, risks and factors that are beyond our control including, without limitation, those more fully described in our filings with the Securities and Exchange Commission ("SEC"), including under the caption "Risk Factors" in our most recent Quarterly Report on Form 10-Q filed with the SEC. Except as required by law, we undertake no duty or obligation to update any forward-looking statements contained in this release as a result of new information, future events, changes in expectations or otherwise.

Contacts

Investors: ir@mongodb.com

Media: press@mongodb.com

 View original content to download multimedia: <https://www.pronewswire.com/news-releases/mongodb-launches-atlas-agent-engine-to-put-ai-agents-in-production-without-a-new-stack-302892753.html>

SOURCE MongoDB, Inc.