



MongoDB Strengthens Foundation for AI Applications with Product Innovations and Expanded Partner Ecosystem

August 11, 2025

New Voyage AI models introduce context awareness and set new accuracy benchmarks—at industry-leading price-performance

MongoDB's AI ecosystem expands AI framework, agentic evaluation, and agentic workflow orchestration capabilities

Approximately 8,000 startups, including Laurel and Mercor, have chosen MongoDB to help build their AI projects

LAS VEGAS, Aug. 11, 2025 /PRNewswire/ -- MongoDB, Inc. (NASDAQ: MDB) today at Ai4 announced a range of product innovations and AI partner ecosystem expansions that make it faster and easier for customers to build accurate, trustworthy, and reliable AI applications at scale. By providing industry-leading embedding models and a fully integrated, AI-ready data platform—and by assembling a world-class ecosystem of AI partners—MongoDB is giving organizations everywhere the tools to deliver reliable, performant, cost-effective AI.



Organizations recognize the business potential of AI. But according to the 2025 Gartner Generative and Agentic AI in Enterprise Applications Survey, 68% of IT leaders felt that they struggled to keep up with the rapid pace at which gen AI tools are rolled out, and 37% agreed that their application vendors drive their enterprise application gen AI strategy. Too many organizations are stuck in the messy middle with their AI implementation, seeing some benefits but not enough to warrant wider adoption.

Businesses express that this gap in AI adoption—a barrier for developers and enterprises alike—is due to the [complexity of the AI stack](#), the importance and challenge of [achieving accuracy](#) for mission-critical applications, and [price-performance concerns](#) that emerge at scale. To address these issues, MongoDB continues to invest in streamlining the AI stack and introducing more performant, more cost-effective models. Customers can integrate Voyage AI's latest embedding and reranking models with their MongoDB database infrastructure. MongoDB has also increased its interoperability with industry-leading AI frameworks—by launching the MongoDB MCP Server to give agents access to tools and data, and by expanding its comprehensive AI partner ecosystem to give developers more choice.

These capabilities are fueling substantial momentum among developers building next-generation AI applications. Enterprise AI adopters like Vonage, LGU+, and The Financial Times—plus approximately 8,000 startups, including the timekeeping startup Laurel, and Mercor, which uses AI to match talent with opportunities—have chosen MongoDB to help build their AI projects in just the past 18 months. Meanwhile, more than 200,000 new developers register for MongoDB Atlas every month.

"Databases are more central than ever to the technology stack in the age of AI. Modern AI applications require a database that combines advanced capabilities—like integrated vector search and best-in-class AI models—to unlock meaningful insights from all forms of data (structure, unstructured), all while streamlining the stack," said Andrew Davidson, SVP of Products at MongoDB. "These systems also demand scalability, security, and flexibility to support production applications as they evolve and as usage grows. By consolidating the AI data stack and by building a cutting-edge AI ecosystem, we're giving developers the tools they need to build and deploy trustworthy, innovative AI solutions faster than ever before."

Accelerating AI innovation with enhanced product capabilities

Voyage AI by MongoDB recently introduced industry-leading embedding models designed to unleash new levels of AI accuracy at a lower cost:

- **Context-aware embeddings for better retrieval:** The new voyage-context-3 model brings a breakthrough in AI accuracy and efficiency. It captures the full document context—no metadata hacks, LLM summaries, or pipeline gymnastics needed—delivering more relevant results and reducing sensitivity to chunk size. It works as a drop-in replacement for standard embeddings in RAG applications.
- **New highs in model performance:** The latest general-purpose models, voyage-3.5 and voyage-3.5-lite, raise the bar on retrieval quality, delivering industry-topping accuracy and price-performance.
- **Instruction-following reranking for improved accuracy:** With rerank-2.5 and rerank-2.5-lite, developers can now guide the reranking process using instructions, unlocking greater retrieval accuracy. These models outperform competitors across a comprehensive set of benchmarks.

MongoDB also recently introduced the MongoDB Model Context Protocol (MCP) Server in public preview. This server standardizes connecting MongoDB deployments directly to popular tools like GitHub CoPilot in Visual Studio Code, Anthropic's Claude, Cursor, and Windsurf—allowing

developers to use natural language to interact with data and manage database operations—and streamlines AI-powered application development on MongoDB, accelerating workflows, boosting productivity, and reducing time to market.

Since launching in public preview, the MongoDB MCP Server has rapidly grown in popularity, with thousands of users building on MongoDB every week. MongoDB has also seen significant interest from large enterprise customers looking to incorporate MCP as part of their agentic application stack.

"Many organizations struggle to scale AI because the models themselves aren't up to the task. They lack the accuracy needed to delight customers, are often complex to fine-tune and integrate, and become too expensive at scale," said Fred Roma, SVP of Engineering at MongoDB. "The quality of your embedding and reranking models is often the difference between a promising prototype and an AI application that delivers meaningful results in production. That's why we've focused on building models that perform better, cost less, and are easier to use—so developers can bring their AI applications into the real world and scale adoption."

"As more enterprises deploy and scale AI applications and agents, the demand for accurate outputs and reduced latency keeps increasing," said Jason Andersen, Vice President and Principal Analyst at Moor Insights and Strategy. "By thoughtfully unifying the AI data stack with integrated advanced vector search and embedding capabilities in their core database platform, MongoDB is taking on these challenges while also reducing complexity for developers."

Expanding the MongoDB AI ecosystem

MongoDB has also expanded its AI partner ecosystem to help customers build and deploy AI applications faster:

- **Enhanced evaluation capabilities:** [Galileo](#), a leading AI reliability and observability platform, is now a member of the MongoDB partner ecosystem, which is designed to give customers flexibility and choice. Galileo enables reliable deployment of AI applications and agents built on MongoDB, with continuous evaluations and monitoring.
- **Resilient, scalable AI applications:** [Temporal](#), a leading open-source Durable Execution platform is now also a member of the MongoDB partner ecosystem. Temporal enables developers to orchestrate reliable AI use cases built on MongoDB, including agents, RAG, and context engineering pipelines that manage and serve dynamic, structured context at runtime. With Temporal's Durable Execution, developers don't need to write plumbing code for resilience or scale. AI applications seamlessly recover across failures, reliably run for a long time, easily handle external interactions, and scale horizontally. Developers can also get visibility into every step of AI workflows to rapidly debug live issues. These partner capabilities significantly expand MongoDB's AI ecosystem for developing AI applications.
- **Streamlined AI workflows:** MongoDB's partnership with [LangChain](#) is redefining how developers build AI applications and agent-based systems by streamlining development and unlocking the value of customers' real-time, proprietary data. Recent advancements include the introduction of GraphRAG with MongoDB Atlas, which enables greater transparency into the retrieval process, fostering trust and providing better explainability of LLMs responses. Another advancement is natural language querying on MongoDB, which allows agentic applications to directly interact with MongoDB data. These integrations empower developers to build reliable, sophisticated AI solutions—from advanced retrieval-augmented generation (RAG) systems to autonomous agents capable of querying data and performing advanced retrieval.

"As organizations bring AI applications and agents into production, accuracy and reliability are of paramount importance," said Vikram Chatterji, CEO and co-founder at Galileo. "By formally joining MongoDB's AI ecosystem, MongoDB and Galileo will now be able to better enable customers to deploy trustworthy AI applications that transform their businesses with less friction."

"Building production-ready agentic AI means enabling systems to survive real-world reliability and scale challenges, consistently and without fail," said Maxim Fateev, CTO at Temporal. "Through our partnership with MongoDB, Temporal empowers developers to orchestrate durable, horizontally scalable AI systems with confidence, ensuring engineering teams build applications their customers can count on."

"As AI agents take on increasingly complex tasks, access to diverse, relevant data becomes essential," said Harrison Chase, CEO & Co-founder at LangChain. "Our integrations with MongoDB, including capabilities like GraphRAG and natural language querying, equip developers with the tools they need to build and deploy complex, future-proofed agentic AI applications grounded in relevant, trustworthy data."

About MongoDB

Headquartered in New York, MongoDB's mission is to empower innovators to create, transform, and disrupt industries with software. MongoDB's unified database platform was built to power the next generation of applications, and MongoDB is the most widely available, globally distributed database on the market. With integrated capabilities for operational data, search, real-time analytics, and AI-powered retrieval, MongoDB helps organizations everywhere move faster, innovate more efficiently, and simplify complex architectures. Millions of developers and more than 50,000 customers across almost every industry—including 70% of the Fortune 100—rely on MongoDB for their most important 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 statements concerning MongoDB's acquisition of Voyage AI. 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: our

customers renewing their subscriptions with us and expanding their usage of software and related services; global political changes; the effects of the ongoing military conflicts between Russia and Ukraine and Israel and Hamas on our business and future operating results; economic downturns and/or the effects of rising interest rates, inflation and volatility in the global economy and financial markets on our business and future operating results; our potential failure to meet publicly announced guidance or other expectations about our business and future operating results; our limited operating history; our history of losses; failure of our platform to satisfy customer demands; the effects of increased competition; our investments in new products and our ability to introduce new features, services or enhancements; social, ethical and security issues relating to the use of new and evolving technologies, such as artificial intelligence, in our offerings or partnerships; our ability to effectively expand our sales and marketing organization; our ability to continue to build and maintain credibility with the developer community; our ability to add new customers or increase sales to our existing customers; our ability to maintain, protect, enforce and enhance our intellectual property; the effects of social, ethical and regulatory issues relating to the use of new and evolving technologies, such as artificial intelligence, in our offerings or partnerships; the growth and expansion of the market for database products and our ability to penetrate that market; our ability to integrate acquired businesses and technologies successfully or achieve the expected benefits of such acquisitions, including the acquisition of Voyage AI; the risk of any unexpected costs or expenses resulting from the acquisition of Voyage AI; the risk of any litigation relating to such acquisition; the risk that such acquisition and the announcement of it could have an adverse effect on our operating results and business generally; our ability to maintain the security of our software and adequately address privacy concerns; our ability to manage our growth effectively and successfully recruit and retain additional highly-qualified personnel; and the price volatility of our common stock. These and other risks and uncertainties are more fully described in our filings with the Securities and Exchange Commission ("SEC"), including under the caption "Risk Factors" in our Quarterly Report on Form 10-Q for the quarter ended October 31, 2024, filed with the SEC on December 10, 2024, and other filings and reports that we may file from time to time 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.

Press Contact:

press@mongodb.com

 View original content to download multimedia: <https://www.prnewswire.com/news-releases/mongodb-strengthens-foundation-for-ai-applications-with-product-innovations-and-expanded-partner-ecosystem-302526003.html>

SOURCE MongoDB, Inc.