

State of AI Agents



Contents

Introduction	03
Key findings	05
Methodology	06
Enterprise AI Agents	07
AI Agent Infrastructure	14
The Agent Ecosystem	19
AI in Production	24
Conclusion	29

Introduction

We're entering a new era of enterprise AI: agentic architectures that produce outcomes accurate for enterprise use cases. These architectures often pair foundation models with enterprise context and tools to independently plan and take action.

Ever since the widespread release of LLMs began disrupting industries three years ago, companies have been searching for ways to harness GenAI's promise of innovation and productivity. Rather than the see-what-sticks approach that many adopted at the outset, enterprises are now honing in on their strategies and driving real operational results. A recent report from [MIT Technology Review](#) reveals that 67% of organizations are already leveraging AI-powered tools, and over half of leaders see agentic AI as a force multiplier for operational performance and decision-making. However, only 19% of organizations have deployed AI agents, and mostly to a limited extent.

AI agents mark the next phase of enterprise AI adoption as organizations transition beyond pilots and chatbots toward agentic systems that can reason and orchestrate workflows. Enterprises that are succeeding are aligning AI use cases to business goals, developing focused agents that complete necessary operational tasks specific to their industry and organization. In the data layer, AI agents require an infrastructure shift that occurs only once every decade.

This report provides a snapshot of how organizations are shaping their data and AI priorities as they push to embed AI across products, services and workflows. Based on insights from more than 20,000 organizations worldwide, including over 60% of the Fortune 500, this report shows how the most innovative organizations are driving their AI agent strategy, transforming architectures to meet the moment and finding success with GenAI.

As the subject of AI agents dominates conversations from boardrooms to breakrooms, this report aims to guide executives and technical professionals in understanding the trends that make organizations successful as they develop their own roadmaps for enterprise AI.

Key findings

Telemetry on the Databricks Platform shows that organizations are thoughtfully using their own data with multiple models and tools to build agentic systems. Investments in governance and evaluations are helping to produce high-quality and accurate outputs that help get AI into production.

Multi-agent systems are becoming the new enterprise operating model

Enterprises are transitioning from single chatbots to multi-agent systems built on domain intelligence.

Use of these systems grew by 327% in just four months.

AI agents are driving core database activities

80% of databases are built by AI agents.

97% of database testing and dev environments are now built by AI agents. This shift is driving the need for a new kind of database called Lakebase.

AI is now part of critical workflows across industries

Most GenAI use cases are focused on automating routine necessary tasks, with 40% related to customer experiences.

Model flexibility is the new AI strategy. 78% of companies are using two or more LLM model families.

AI evaluations and governance are the building blocks of production

AI evaluations are critical to ensuring high-quality outputs. Companies that use evaluation tools get nearly 6x more AI projects into production.

Companies using AI governance put over 12x more AI projects into production. AI governance is a top investment priority, and grew 7x in nine months.

Methodology:

How did Databricks create this report?

The 2026 State of AI Agents is built from fully aggregated, anonymized data collected from our customers based on how they are using the Databricks Data Intelligence Platform.


This report focuses on trends in AI and agent strategies, use cases, and infrastructure. The customers in this report represent every major industry and range in size from startups to many of the world's largest enterprises. Whenever possible, this report presents and analyzes data from November 1, 2024, to October 31, 2025. Due to the fast pace of AI development, new products are constantly launched, and in many cases, a full year of data is not available. These charts are noted accordingly.

Usage is measured by the number of active, engaged customers. Minimum bars are used when needed to ensure growth rates are accurately represented. When possible, we provide year-over-year (YoY) comparisons to showcase growth trends over time.

Enterprise AI Agents

The meteoric rise of AI frontier models has laid the groundwork for today's agentic AI systems. While AI chatbots were a foundational early application and continue to be a prevailing use case, today's innovators are turning their attention toward developing multi-agent systems that can autonomously plan and carry out complex and specialized tasks.

Usage of Agent Bricks



Databricks Agent Bricks use cases:

- **Custom LLMs:** generate text-based results for domain-specific tasks
- **Supervisor Agent:** orchestrates agents and tools to execute complex tasks
- **Knowledge Assistant:** custom-trains Q&A chatbots on enterprise documents
- **Information Extraction:** transforms unlabeled text documents into structured tables

Note: Data range is June through October 2025 due to June launch of Databricks Agent Bricks.

The next frontier of enterprise AI is multi-agent systems

To understand how organizations are actually using AI agents within their own workflows, we look at the usage of the four types of agents available in Databricks Agent Bricks.

327% growth in multi-agent workflows

Enterprises are transitioning from single chatbots to multi-agent systems that autonomously orchestrate and execute full workflows.

Supervisor Agent creates systems comprised of multiple agents that work together to complete tasks across specialized domains. These agents are auto-optimized on an organization's own data and through various optimization techniques.

A Financial Services organization may build a multi-agent system to handle intent detection, document retrieval and compliance checks, delivering personalized responses to clients. Since Databricks launched Supervisor Agent in July 2025, it has quickly become the No. 1 agent use case and accounted for 37% of all Agent Bricks usage in October 2025.


Information Extraction represents 31% of usage

The second most commonly used agent, Information Extraction, reflects companies' need to leverage both structured and unstructured data. By converting unlabeled text into structured tables, Information Extraction enables organizations to tap into all their data for their AI initiatives. For example, retail organizations can easily pull product details, prices and descriptions from complex supplier PDFs, even if the documents are formatted differently.

Tech builds nearly 4x more multi-agent systems than any other industry

Tech companies lead the way in adopting these agentic systems, pointing to their ability to break down business challenges into defined problems that can be solved by specialized, coordinated AI agents.

Supervisor Agent Usage, by Industry




Note: Data range is July through October 2025.


Top AI use cases: Companies automate routine tasks

Our data shows that enterprises are taking a pragmatic approach with their AI initiatives by solving real-world business problems.


Top AI Use Cases



Top AI Use Cases by Industry



Top AI Use Cases by Region



Use Case Categories ↓

- Market Intelligence & Strategic Analytics
- Business Operations & Automation
- Customer Experience & Engagement
- Healthcare & Scientific Knowledge Management

Among our categorized AI use cases, the top 15 focus on automating necessary, but routine, daily tasks. This trend stays consistent across all industries and global regions, though use cases are often tailored to an industry's unique operational needs. For example, the top use case in Healthcare & Life Sciences focuses on medical literature synthesis, while Manufacturing & Automotive and Energy & Utilities prioritize predictive maintenance.

Automating routine tasks can build rigor and confidence, providing the organizational base for agent systems. And AI-driven automation isn't just for technical experts. Many of these tasks can be executed by applying AI directly onto proprietary data and data pipelines using built-in functions.

40% of the top AI use cases focus on customer experience and engagement

According to a study by Deloitte Digital, nine in 10 customer experience leaders are confident that AI has the potential to improve the customer experience. Our data shows companies are now using AI with that intent: 40% of AI use cases focus on practical customer concerns such as customer support, customer advocacy, onboarding and personalized marketing content.

AI Agent Infrastructure

AI agents are fundamentally changing infrastructure as they've gone from barely a blip in database management to driving core data work functions. Together with the surge in "vibe coding," this has profound implications for database requirements.


Agents are at the helm of database operations

AI agents now create 80% of databases

Gartner estimates that by 2028, enterprises will use vibe coding techniques and tools to create 40% of new production software. As vibe coding fuels the explosion of data intelligent applications, databases must scale just as fast.

Agents make this scale possible, and we've seen this shift already happening within our own data. Telemetry in Neon, a serverless Postgres database acquired by Databricks and the core technology behind Databricks Lakebase, showed that the number of databases created by AI agents increased from 0.1% to 80% in just two years.

Neon Components Created by AI Agents



AI agents create 97% of database branches

AI agents have nearly taken over the creation of database branches, slashing the time required to clone databases from hours to seconds.

Branches describe isolated variants of a database that are fully functional and version-controlled. Agents quickly create these ephemeral environments as a place to test, experiment and develop safely, and even rewind to an earlier state without the need for heavyweight provisioning. Two years ago, 0.1% of database branches were created by AI agents; today that number is 97%.

Agents can spin up branches, sandboxes and experimental pipelines continuously, triggering an exponential volume of platform operations at a pace far beyond what humans can keep up with.

This speed of development is critical to scale with the explosion of vibe-coded applications. When building is democratized, organizations can quickly prototype and create tailored data intelligent applications of all sizes and for specific use cases, which puts increasing pressure on the operation layer. Agents help database teams scale with this new era of app development.

Databases for the AI era

Online transaction processing (OLTP) databases are designed to efficiently handle high volumes of real-time transactional data. Traditional OLTP databases have powered decades of business operations, but they were built around human workflows with predictable transactions, infrequent schema changes and manual provisioning.

As enterprises adopt agentic AI, these assumptions no longer hold. Agents generate continuous, high-frequency read and write patterns and need to create and tear down environments programmatically. They perform complex, read-intensive queries as part of multistep reasoning loops. This new behavior places demands on databases that traditional OLTP systems were never architected to support. These agentic workflows require infrastructure that can support massive concurrency,

millisecond-scale branching, isolation and real-time responsiveness at a scale beyond manual human workloads.

This shift has given rise to a new category of operational databases built for the AI era, systems that combine transactional performance with the elasticity, programmability and scale needed for autonomous agents.

Lakebase exemplifies this new class of operational database, enabling AI agents to create, branch and orchestrate Postgres transactional databases in milliseconds. By running OLTP workloads on the same object store that is the foundation of the Databricks lakehouse, Lakebase delivers the low latency and high-volume operations required for data intelligent applications and AI agent systems.


The rise of the citizen AI app developer

The rise of “vibe coding,” where users describe what they want in natural language and let AI generate the code, means citizen AI app developers can now build working prototypes without deep technical expertise. Today, even business users can create apps with minimal coding skills, leading to the democratization of AI throughout the organization.

The citizen AI app developer is already in motion. Over 50K data and AI apps have been created since the Public Preview of Databricks Apps, with a 250% growth rate over the past six months.

These proofs of concept then become the starting point for collaboration: data teams and application developers can take what works, refine it and productionize it as an AI app ready for enterprise deployment.

Databricks Apps Created




The Agent Ecosystem

Using multiple model families not only helps companies align their use cases with the most performative LLMs, but also future-proof against vendor lock-in. An open, unified platform supporting all models is key to this strategy.

Multi-model strategy takes off


As of October 2025, 78% of companies use two or more LLM model families, such as GPT, Claude, Llama, Gemini and Qwen. The share of companies using three or more LLM model families rose sharply, from 36% as of July 2025 to 59% as of October 2025. Having a platform that can easily integrate any model from any AI provider, whether proprietary or open source, and combine multiple models within a single agent system, will unlock the widest set of use cases.

Number of LLM Model Families Used per Company




Multi-model AI varies across industries

As multi-model AI trends toward widespread adoption, certain industries are setting the pace. Companies in the Retail industry are the most likely to use multiple LLM model families, with 83% of companies using two or more model families. AI is transforming the Retail industry across all angles, from productivity to unlocking new ways to engage customers. Leveraging multiple LLM families is a strategy for retail companies to balance performance and cost tailored to the specific tasks.



Real-time serving inference outpaces batch processing

Our data shows that 96% of all requests are processed in real time. Applications might include interactive AI experiences such as copilots, customer support assistants and personalization. Real-time requests are especially crucial in industries where subsecond inferences have a significant impact on outcomes.




Note: Data range is May through October 2025.

Industry and regional trends of real-time serving

The Tech industry has the highest ratio of real-time requests, processing 32 real-time requests for every one batch request.

The Healthcare & Life Sciences industry, which has the second-highest ratio of real-time inferences (13 requests for every one batch request), often addresses critical situations such as emergencies, continuous patient monitoring and diagnoses, and personalized treatments.



Note: Data range is May through October 2025.

AI in Production


Achieving value from AI projects requires getting them into the real world. But in many organizations, there's still a large gap between experimentation and production. According to a [2025 MIT NANDA](#) report, 95% of generative AI pilots fail to get into production. Across industries, executives want to know: What's the key to productionizing AI and agents?

7x investments in AI governance and security

In opening the door to powerful new possibilities, agentic AI also ushers in high technical complexity and autonomy, as well as new inherent risks.

Unified governance helps enterprises ensure their AI development aligns with business goals, meets security and legal obligations, and accounts for regulatory risks. Use of AI Gateway, Databricks' AI governance solution, has skyrocketed 7x since January 2025 as companies race to implement AI solutions and trigger a proliferation of data and AI apps, models and agents.

Usage of Databricks AI Products



The operational guardrails of agentic AI

A 2024 global survey conducted by Economist Impact found that 40% of respondents believed their organization's AI governance program was insufficient in ensuring the safety and compliance of their AI assets and use cases.

Our data shows that governance is the precondition that enables enterprises to deploy agentic systems at scale. Unified governance makes agentic AI possible by serving as a layer that defines how data is used, setting guardrails and rate limits. It also establishes structured accountability within organizations, helping to ensure that AI decisions are aligned with evolving ethical standards.

Companies that actively use AI governance put 12x more AI projects into production

12x

Evaluations: The powerful catalyst to produce high-quality outputs required for AI in production

AI evaluation tools are frameworks designed to systematically measure, test and improve the quality and reliability of AI models at all stages of deployment.

Governance and evaluations have a symbiotic relationship in the development of AI models. While governance provides the guardrails and control panel for agents, evaluations monitor and measure agent behavior throughout their lifecycle, enabling governance to adapt in real time as agents learn or as environments change.

AI evaluation has become a powerful lever for deploying agentic systems at scale. Evaluation goes beyond general benchmarks to create custom benchmarks specific to an enterprise's data and the goal task. The agent then continuously tests accuracy, safety, fairness and compliance and seeks to improve the outputs. The rapid feedback from evaluations empowers teams to detect problems early and iterate strategically until they are confident about moving trustworthy, high-quality models into production.

The MIT Technology Review Insights survey found that only 2% of respondents rate their organization's AI performance highly in terms of delivering measurable business results.

Domain-specific evaluations validate knowledge and decisions grounded in enterprise data, allowing teams to tie evaluation metrics to business KPIs (such as CSAT, handle time and revenue lift), making improvements actionable. This continuous evaluation transforms AI agents from static tools into learning systems that improve over time and have the potential to unlock scalable business impact.

Companies actively using evaluation tools get nearly 6x more AI projects into production

6x

Conclusion

As enterprises adopt agentic AI, old rules and processes are being turned on their heads.

Agents have begun running the data layer in a massive operational shift, and enterprise AI now includes coordinated and autonomous workflows based on domain intelligence.

So what does this mean for companies today? For executives and data professionals, the challenge now is not about selecting the right model or agent use case, but how to effectively use agents with enterprise context to produce high-quality, accurate outputs. Our findings show leaders in enterprise AI are:

- Choosing the right model and/or multi-agent system to solve specific problems
- Using AI to support high-volume mundane tasks
- Investing in unified governance and evaluation tools, which have proven to be key drivers to getting AI into production

The takeaway

As agentic AI redefines what's possible, the boldest organizations will set the pace of transformation.

Contributors

State of AI Agents is the result of the close collaboration between many teams and domain experts. Thank you to the following Bricksters, whose contributions and expertise shaped the insights shared in this report.

Data Leaders

- Kunal Marwah
- Mason Force
- Ram Murali

Technical Council

- Ahmed Bilal
- Avesh Singh
- Alkis Polyzotis
- Corey Zumar
- Elise Gonzales
- Emīls Pauls Eglītis
- Jake Brutlag
- Jeffrey Gu
- Vibhu Saxena

Content Lead

- Mikaila Garfinkel

Art Director

- Jarod Octon

