

Connecting AI agents to enterprise knowledge



Preface

“Connecting AI agents to enterprise knowledge” is an MIT Technology Review Insights report sponsored by Neo4j. This report, which is based on survey research and executive interviews, seeks to understand how organizations are working to imbue their AI agents with the knowledge needed to make good decisions. Denis McCauley was the author of the report, Virginia Wilson was the editor, and Nicola Crepaldi was the publisher. The research is editorially independent and the views expressed are those of MIT Technology Review Insights.

We would like to thank the following executives for their time and insights:

Sudhir Hasbe, President and Chief Product Officer, Neo4j

Jim Jacob, Executive Director of Digitalization, Advanced Analytics, and Artificial Intelligence, Cummins

Thomas Luu, Director, Anti-Counterfeiting Team, Gilead Sciences

Sanjeev Mohan, Principal, SanjMo

Methodology

In June and July 2026, MIT Technology Review Insights conducted a global survey of 300 chief data and analytics officers, chief information officers, chief technology officers, chief AI officers, and other senior data and technology executives. The survey respondents work in organizations across nine industries. All of their organizations earn \$500 million or more in annual revenue, with 38% earning over \$10 billion annually.

In addition to the survey, a series of in-depth interviews with senior technology executives offer firsthand insights into the readiness of organizations' data systems to support agentic AI.

CONTENTS

01 Executive summary	4
02 Underperforming agents.....	5
Points of project failure.....	6
No success without the right foundations.....	7
An agent that works for analysts and senior management alike.....	8
03 Understanding the knowledge gaps.....	9
The knowledge imperative.....	10
Seeking a semantic edge	10
The components of knowledge for AI.....	11
04 Toward context-rich agents.....	12
Next step: Adding a knowledge layer	13
How Gilead Sciences combats fraud with a knowledge layer	14
Tools for the challenge	14
05 Conclusion	16

01

Executive summary

For all the data that AI systems continually amass and analyze, enterprise AI agents often suffer from a curious shortcoming: a lack of knowledge. More than data, knowledge is the understanding of what the data means in the context of individual organizations. AI agents need this understanding to reason about situations, make decisions, and ultimately take actions. Without sufficient knowledge, agents are prone to making flawed and unreliable decisions.

A lack of knowledge, our research finds, is a major reason agentic AI use cases never make it to production. Competitive pressure is making it urgent to address this. Organizations need to deploy and scale more of their agentic projects to capture the efficiency gains AI promises. Falling short risks wasting the investment already sunk into these projects, and it cedes ground to rivals already putting their agents to work more effectively.

The purpose of this report, which is based on a survey of 300 data, AI, and other technology executives, is threefold. First, it seeks to gauge organizations' agentic knowledge capabilities (i.e., their ability to give AI agents a full contextual understanding of the data they ingest) across semantic knowledge, episodic memory, and procedural knowledge. Second, the report probes the challenges organizations face in improving access to knowledge and ultimately to getting more agent use cases into production. Third, it explores the measures organizations are taking to overcome these challenges.

The key findings include the following:

Data and knowledge weaknesses consistently stall AI agent progress. On average, only around a third (34%) of organizations' agentic AI projects make it into production. Even high-tech firms struggle with this. Legacy data systems, security and privacy concerns, and a lack of knowledge and context are the key points of failure.

Strong knowledge capabilities correlate with agent success. A small group of production leaders (organizations where on average 61% of agentic projects advance beyond pilot) have stronger knowledge capabilities than the rest, especially when it comes to semantics. This advantage tracks closely with their higher production rate.

Fragmented data hugely complicates knowledge access. Data fragmentation (the inadequate sharing of data across systems) was most commonly cited as a top challenge to expanding agents' access to knowledge (cited by 55%). Production leaders, by contrast, are more likely to see security and privacy concerns as a major concern (cited by 72% of this group).

Most firms aim to strengthen the link between data and agents. Among steps that can yield higher quality agent decisions, executives expect the biggest impact to come from strengthening the structural foundation between the organization's data and its AI agents. The experts we interviewed see a knowledge layer as a prime way to achieve this.

Investment priorities to boost knowledge range from pipelines to knowledge graphs. To expand agent access to knowledge, organizations will prioritize investments in retrieval technologies, like ingestion pipelines, AI-ready APIs, and retrieval-augmented generation (RAG); in AI evaluation agents; and in knowledge graphs.

02 Underperforming agents

In the two years since agentic AI became available for enterprise use, AI agents have found a home in many organizations, but it is not yet a comfortable one. Gartner predicted last year that over 40% of current agentic AI projects would be canceled by the end of 2027.¹ Our research suggests that the figure may turn out to be higher.

Just over half (57%) of the executives we surveyed say their organizations use AI agents widely today. That includes widescale experimentation and piloting, as well as deploying in production. But most organizations are having difficulty advancing their use cases into production. The average success rate across the group is just 34% (see Figure 1).



Sanjeev Mohan, principal of advisory firm SanjMo, believes that even the 34% figure may be an overestimate. “I think it’s closer to the single digits,” he says. Jim Jacob, executive director for digitalization, advanced analytics, and artificial intelligence at Cummins, holds a similar view: “For companies needing 100% accuracy of results, the figure is more likely to be around 20% or less. The higher figure is probably more for companies with looser accuracy tolerances.”


Perhaps unsurprisingly, agent use is most advanced among the survey’s technology companies (including IT, software, and other technology providers). But even in this sector, just 42% of executives say their agentic projects make it to production.

Figure 1: Only around one-third of agentic AI projects advance beyond the pilot stage



Source: MIT Technology Review Insights survey, 2026

Figure 2: An average 61% of production leaders' agentic AI projects reach production



Source: MIT Technology Review Insights survey, 2026

Project success rates are considerably lower elsewhere, and they are lowest of all in consumer packaged goods (CPG) and in healthcare and life sciences (see Figure 1).

A select group (13%) of organizations, which we call production leaders, are successfully deploying over half their agentic AI projects (61% on average). The remaining respondents – the production laggards – are deploying half or fewer of their agentic AI projects (30% on average) (see Figure 2).

The production leaders exhibit some distinct characteristics. For example, their AI agents have stronger knowledge capabilities than the laggards. These leaders are also less likely than the rest to struggle with legacy data infrastructure and fragmented data.

For organizations with high project failure rates, the costs are likely to be keenly felt. They take the form of wasted investments – for example, in funds allocated, engineering time, and the cost of data specialists. They also involve lost opportunities, in areas like productivity gains, customer experience improvements, and new revenue generation. High failure rates can also dent the trust of users and management, making it harder to build a case for future agentic AI investments.

Points of project failure

Why do so few agentic projects survive piloting? An AI director in the entertainment industry emphasizes the scale of the challenges that such projects must overcome before they can advance to production: “In my experience, many AI pilots amount to impressive demonstrations that create a lot of ‘aha’ moments,” he says. “Production, on the other hand, requires proving that the complete system can operate repeatedly, securely, economically, and under changing conditions. Not every pilot takes this into account. Pilots may succeed on curated examples, but most don’t amid real-world variability.”

“As with all types of AI, too many agentic use cases fail because teams have gone after something speculative that doesn’t make much business sense.”

Sudhir Hasbe, President and Chief Product Officer, Neo4j

Figure 3: Legacy systems, security concerns, and insufficient agent knowledge are the top reasons that agentic AI projects do not reach production



Source: MIT Technology Review Insights survey, 2026

Unclear business value is the most fundamental weakness for many projects, says Sudhir Hasbe, president and chief product officer of Neo4j. “As with all types of AI, too many agentic use cases fail because teams have gone after something speculative that doesn’t make much business sense,” he explains.

Technology factors also contribute to project failure. For the survey respondents, three such factors stand out: problematic agent integration with legacy data systems (cited by 55%), concerns with security and/or data privacy (53%), and insufficient agent knowledge and context (46%). Inadequate auditability of agent decisions (44%) and agent hallucination (42%) are also common points of failure (see Figure 3).

No success without the right foundations

The main reason why so few agentic projects succeed is that organizations do not pay enough attention to putting the foundations in place, says Jim Jacob, executive director of digitalization, advanced analytics, and artificial intelligence at Cummins. “Companies expect agents to be able to do things that were never done before,” he says. “But the fundamentals need to be in place for that to become a reality. Your data must be accurate, your connections must be reliable and accurate, and the way data gets orchestrated is essential.”

When the connections between sourcing, manufacturing, logistics, and planning are well orchestrated, says Jacob, an agent can indeed meet high expectations. “The agent can say, ‘You’ve asked me for a forecast for raw material X, but by the way,

widget Y, which is highly dependent on raw material X’s success, is in short supply. So, procure widget Y first, and then buy raw material X.’” For Jacob, that “by the way” is the most important intelligence that agentic AI delivers. It can provide far deeper solutions to what was asked.

“Industrial enterprises are living organisms in which products, R&D, external timing challenges, management, mergers and acquisitions, and divestitures are all connected,” says Jacob. “For such realities, connectivity, ontologies, and knowledge maps must be adaptable to the changing environment. Gone are the days of static tables and reports as decision foundations. AI and knowledge maps help build that adaptability. Investing time getting that foundation right is a very important step.”

An agent that works for analysts and senior management alike

"If it's to succeed, every pilot needs a clearly measurable business outcome," says an AI director in the entertainment industry. It also needs metrics to chart progress against that outcome. Drawing from his experience with enterprise agentic AI initiatives, the AI director explains how these types of enterprise agents can work.

In his experience, these agents tend to perform better when connected to curated and well-contextualized sources. "This makes the reasoning traceable and creates clear escalation paths depending on the risk profile of the action it's

suggesting. This determines whether it can make an autonomous decision or if a human needs to be brought into the loop."

The agent's success is measured based on its adoption and repeat use. "We look at measures such as adoption and repeat usage, reduction in search or analysis time, accuracy against reviewed test sets, escalation rates, user corrections, and completion rates," the AI director says. "Together, these types of measures provide a much better indication of production value than a successful demonstration alone."


Respondents in the healthcare, life sciences, and CPG sectors most commonly cite insufficient knowledge and context as a challenge to their agentic AI projects. In healthcare, it is the most likely point of failure of all (cited by 61%). According to Hasbe, an overly simplistic approach to data storage and retrieval also contributes to project failures, which can exacerbate a lack of agent knowledge and context.

The production leaders are far less likely to ascribe project failures to knowledge and context limitations (cited by just 28%) or legacy data systems (41%). For this group, security and privacy concerns and hallucination do more to stall projects. "Given recent stories about agents jumping the sandbox and launching attacks, executives are understandably wary about approving autonomous agents for production and potentially putting their jobs on the line," says Mohan.

"Companies expect agents to be able to do things that were never done before. But the fundamentals need to be in place for that to become a reality. Your data must be accurate, your connections must be reliable and accurate, and the way data gets orchestrated is essential."

Jim Jacob, Executive Director of Digitalization, Advanced Analytics, and Artificial Intelligence, Cummins

03 Understanding the knowledge gaps



If undetected, technology challenges that limit agentic AI rollout – including problems integrating legacy data systems, security and data concerns, and insufficient knowledge and context – can manifest themselves in incorrect or unreliable agent decisions. This can lead to damaging real-world consequences, as demonstrated in the following examples:

- A coding agent erroneously decided to delete a software company's entire production database. This meant that the car rental companies that use the software were unable to manage reservations and vehicle assignments for their customers for more than a day.²
- A health insurer's agent overrode physician recommendations and denied post-acute care to elderly patients. An investigation found that the model failed to account for comorbidities and individual patient complexity that physicians recognized as critical to recovery.³

- A fast-food restaurant's flawed chatbot led to massive over-ordering by a customer, causing reputation damage when news of it went viral. The system failed to identify an obviously erroneous order volume, and it also lacked adequate knowledge about a product the company was selling.⁴

Weaknesses in data systems – in infrastructure, architecture, or the underlying data – are likely to have played some role in these failures. Missing or error-ridden data is a common challenge to deploying agentic AI projects, according to 38% of survey respondents (see Figure 3).

Before launching agent pilots, data cleanliness is vital, says Thomas Luu, director of the anti-counterfeiting team at biopharma company Gilead Sciences. "For an agent to provide good outputs, it's critical that the information and data that it needs is very clean and crisp," he says. "We did a lot of upfront work cleaning the data and centralizing it before giving our fraud detection agent access to it."

"In fraud detection, knowledge gives us the context to help understand not only what's going on currently, but also how an entity has evolved their business transactions historically. That will reveal signals of fraud, waste, and abuse within the current data. That's important because we're trying to detect fraud proactively."

Thomas Luu, Director, Anti-Counterfeiting Team, Gilead Sciences

The knowledge imperative

Luu also stresses the importance of imbuing agents with knowledge. “In fraud detection,” he says, “knowledge gives us the context to help understand not only what’s going on currently, but also how an entity has evolved their business transactions historically. That will reveal signals of fraud, waste, and abuse within the current data. That’s important because we’re trying to detect fraud proactively.”

But expanding agent access to the organization’s knowledge is easier said than done. “The fundamental issue is that the large-scale data systems that so many organizations built over the last 30 years were designed for data consumption by humans, not by machines,” explains Hasbe.

When asked about the challenges of expanding agent access to knowledge, the surveyed executives cited factors similar to those in overall agentic project failure. At the top of their list, mentioned by 55%, is data fragmentation – in other words, the inadequate sharing of data across systems (see Figure 4). “This is a big problem,” says Hasbe. “You can’t expect an agent to reason if your data is siloed in 15 systems and across disconnected tables.”


Another major challenge for respondents are security and privacy concerns (cited by 48%). These are by far the toughest challenges for financial industry organizations (74%) and those in the tech sector (62%). Our group of production leaders also put this at the top of their list.

“Data fragmentation and security are fundamental challenges, and they interact with each other,” says the AI director in the entertainment industry. “Knowledge is distributed across structured systems, documents, conversations, and human expertise. When organizations try to centralize or index all of it, it can create a new security problem. It is a difficult challenge to balance retrieval from fragmented sources without violating underlying security policies.”

Seeking a semantic edge

The survey respondents are broadly positive about the knowledge capabilities of their AI agents today. But that assessment comes with caveats. Four in five (80%) describe the semantic knowledge their agents have access to as strong. But just a quarter (26%) say it is very strong, with 54% describing it as somewhat strong. Just 10% say their agents’ episodic memory capabilities are very strong, while 47% deem them somewhat strong. The analogous figures for procedural knowledge capabilities are 31% and 38%, respectively.

Figure 4: Fragmented data along with security and privacy concerns limit the knowledge organizations are able to provide to agents



Source: MIT Technology Review Insights survey, 2026

The components of knowledge for AI

An enterprise's knowledge consists of several elements, but the most fundamental for AI are semantic knowledge, episodic memory, and procedural knowledge.

- **Semantic knowledge** refers to the organization's repository of facts, rules, and concepts. It is a key determinant of context, helping an agent understand what the organization's most important relationships and metrics are. When an agent operates without semantic knowledge, it can misinterpret user questions, generating inaccurate responses and eroding trust.
- **Episodic memory** is experience-based knowledge accumulated from previous interactions and agent decisions. "Memory is critical," says Hasbe. "As an agent performs tasks over a period of time, and especially if they're long-running tasks, memory helps it remember what it did before and what it's supposed to do in the future."
- **Procedural knowledge** tells an agent how the organization executes different actions. It gives an agent the knowledge to, for example, debug code in different situations, place an order with a supplier, or book a flight for a staff member.


The production leaders are much more confident of their agents' knowledge capabilities, especially when it comes to semantics. This is manifested most clearly when it comes to semantic knowledge, which 62% of this group describe as very strong. It is this advantage that most sets the production leaders apart from the rest of the surveyed organizations and helps explain why the leaders are so much more successful at getting their agentic projects into production.

For the production laggards, the contrast between their low project success rate and the broadly positive self-assessment of their semantic knowledge suggests that they may be overestimating the latter capability.

"The fundamental issue is that the large-scale data systems that so many organizations built over the last 30 years were designed for data consumption by humans, not by machines."

Sudhir Hasbe, President and Chief Product Officer, Neo4j

Figure 5: Production leaders outshine the rest in semantic and other knowledge capabilities



Toward context-rich agents



Organizations can improve AI agent outcomes through a combination of data architecture, organizational, and human initiatives. One such human-focused effort is the upskilling or reskilling of data and AI specialists, which 36% of respondents (the second highest proportion) cite as having a positive impact on the quality of AI agent decisions (see Figure 5). Strengthening human monitoring of agent decision-making was also selected by 23% of respondents. “The importance of human oversight cannot be overstated,” says Luu. “It is a key factor in an agent’s success.”

According to the AI director in the entertainment industry, humans should not be inserted into every step of the agent workflow by default. “That adds work and destroys efficiency gains. However, oversight remains essential because strong pilot performance doesn’t guarantee production reliability without human checks,” he says.

At Cummins, the need for human oversight is dictated partly by the need for decision accuracy in any given process, says Jacob. “When you order, you put


money on the table and carry inventory, so you may only be able to tolerate a very low variance,” he says. “The margin of error could be slightly higher when forecasting. But in financial reporting to the Street, for instance, you can’t have any errors; accuracy is required. That’s when you need to fine-tune the agent down to a repeatable process. We tend to keep a human in the loop for four, five, or six cycles to make sure the agent output and human output are very close, shrinking that delta to the minimum possible before running the agent autonomously – but always with human-in-the-loop for final approvals.”

A third of respondents (34%) agree that updating their data architecture will have a favorable impact on decision quality (the third most popular initiative overall) (see Figure 6). But the single most beneficial initiative will be strengthening the structural foundation between the organization’s data and its AI agents. More than four in 10 (44%) expect this to have the largest and most positive impact on agent decision quality. Among sectors, this view is held most strongly in retail (54%), automotive (53%), and CPG (48%).

“When you order, you put money on the table and carry inventory, so you may only be able to tolerate a very low variance. The margin of error could be slightly higher when forecasting. But in financial reporting to the Street, you can’t have any errors.”

Jim Jacob, Executive Director of Digitalization, Advanced Analytics, and Artificial Intelligence, Cummins

Figure 6: Strengthening the link between the organization's data and its AI agents will boost the quality of agent decisions



Source: MIT Technology Review Insights survey, 2026

For those who emphasize the importance of bolstering that structural connection, the anticipated business benefits are tangible. The principal outcomes they expect from this over the next two years are greater operational efficiency (cited by 57%), improved productivity (48%), and reduced risk from incorrect agent decisions (41%) (see Figure 7).

Next step: Adding a knowledge layer

Building a knowledge layer into the data architecture can strengthen that structural foundation, according to several of the experts we interviewed. Mohan explains the role of the knowledge layer: "It bridges the gap between the business request and the technical layer, and ensures that terms are crystal clear and

consistently defined. Take the question, 'Did this product make a profit in the Midwest in Q1 of fiscal 2026?' The knowledge layer tells the agent what's fiscal, what's Q, what's region, and what's profit."

According to Jacob, the knowledge layer can help industrial companies like his overcome the data fragmentation challenges that often hinder the effectiveness of AI agents. "Most firms in our industry are digital non-native, and disparate data systems are rampant in them," he says. "Knowledge maps create a front-end view of structured common data. The connections go into disparate datasets, but they create a common dataset view that an agent can depend on."

Figure 7: Building a strong link between data and agents could bring gains that impact the bottom line



Source: MIT Technology Review Insights survey, 2026

How Gilead Sciences combats fraud with a knowledge layer

According to Thomas Luu, director of the anti-counterfeiting team at Gilead Sciences, the company's data platform and knowledge layer have helped it move beyond anomaly detection to uncover extended networks of medical and pharmaceutical fraud.

When Luu and his fraud team started on their agentic AI journey, they could pick up things that looked out of the ordinary from deviations in the data. "But the most clever fraudsters are the ones hiding in the averages," he says. "Their behavior is not easy to pick up, because it often looks normal on paper."

Incorporating a knowledge layer has helped the team to bring those entities to light, says Luu. He offers an example of how they use it. "We do a lot of investigating of open source intelligence (OSINT) to find context about particular entities, and we add that data to the knowledge graph – a core component of the knowledge layer. Our agent may then determine that an entity is one or two steps (or 'hops' in the graph world) away from the actual

fraud that's happening. They may not be directly committing the fraud, but they're close to it."

Seemingly legitimate pharmacies often push their fraudulent activity out to burner pharmacies, according to Luu. "This might be a storage facility or an individual pharmacist the main organization has hired. When we detect something anomalous or fraudulent, we'll shut down the burner pharmacy, but another will always pop up. The main pharmacy, meanwhile, continues receiving the illegitimate product while still operating within normal thresholds."

With its own historical data and OSINT added into the knowledge graph, the team's agents can now help the company identify entire networks of fraudulent operation, says Luu. "The knowledge layer has enabled us to identify all the main culprits and understand how the schemes work and where they are centralized. That's given us a big competitive edge in combating fraud."

"Without the contextual information provided by the knowledge, all we have are statistical signals," says Luu. "In our industry, the fraudsters know the system. They know how medical and pharmaceutical billing work. They know how to take advantage of government and manufacturer rebates. Because that data always lived in silos, it's historically been tough to detect fraud. But when you look at all the data together, which we can do now with our knowledge layer, it becomes obvious that there's something going on."

Tools for the challenge

Improving agent access to organizational knowledge, and ultimately improving the quality of agent decisions, naturally requires some upgrades to data infrastructure and tools. "To deliver greater knowledge and context, current data estates need better data retrieval pipelines, upgraded data APIs, and also knowledge graphs," says the AI director in the entertainment industry. "The goal isn't just connecting agents to more data but delivering the right context with the right permissions for a specific business outcome."

"The knowledge layer bridges the gap between the business request and the technical layer, and ensures that terms are crystal clear and consistently defined."

Sanjeev Mohan, Principal, SanjMo

“Previously, we couldn’t look at a lot of datasets together because they don’t talk to each other. With graph databases, we’ve been able to make those datasets communicate with each other.”

Thomas Luu, Director, Anti-Counterfeiting Team, Gilead Sciences

These technologies all feature in the list of data system investments that the surveyed firms will prioritize over the next two years. Automated ingestion pipelines (cited by 54% overall), AI-ready APIs (50%), and RAG (37%) will reduce latency and ensure that agent knowledge is kept up to date (see Figure 8).


AI evaluation agents (selected by 37%), meanwhile, will augment the work of human analysts in identifying agents’ knowledge gaps and detecting the retrieval of irrelevant or incorrect data, for example.

And knowledge graphs (35%) will do the job of organizing data as a network of interconnected entities and relationships. The knowledge graph provides the contextual understanding agents need to reason accurately over data and thus is a critical component

of the knowledge layer. Importantly, it can do this across numerous disparate datasets. “Previously, we couldn’t look at a lot of datasets together because they don’t talk to each other,” says Luu. “With graph databases, we’ve been able to make those datasets communicate with each other.”


Jacob cautions against the risk of driving up costs when creating a knowledge layer. “You can spend huge amounts of money on tokens, memory, and compute,” he says. “Our approach is to build a ‘digital core’, which uses graph connections to link existing datasets and establish context in a single location. Each use creates digital neural pathways locally without driving up cloud, compute, or token costs. If you rely solely on direct large language model translation without this layer, you will consume immense enterprise resources.”

Figure 8: Ingestion pipelines and AI-ready APIs are priority investment targets to expand agent access to knowledge



05

Conclusion



This study has highlighted the link between organizational knowledge and successful AI agent use cases. It has also shed light on the data system weaknesses that keep agents from leveraging that knowledge. Several of the experts we interviewed affirm that integrating a knowledge layer into the data architecture will go a long way toward addressing those weaknesses.

It remains to summarize the bigger rewards to be gained from strengthening the foundation between the organization's data and AI agents in this way. Our research points to three interrelated areas in which the business as a whole is likely to benefit from tying enterprise data and agents more tightly together:

Trust in AI: The better an agent's understanding of what's important to the business, the less chance of hallucination, irrelevant results, or incorrect decisions. The more reliable the performance of agents, the greater the trust they will engender across the organization and the faster adoption will grow.

Transparency of decision-making: When agent decisions are grounded in data that fully takes account of the organization's policies and objectives, and does not infer them from other sources, traceability of those decisions can be maximized. Trust and compliance both benefit.

Unlocking scale and autonomy: Structured enterprise knowledge will facilitate the reuse of context across agents, which will accelerate time to production and scaling to more parts of the business. And while humans will stay close to agents for some time to come, a trusted knowledge foundation will make it easier for organizations to increase the scope of agentic automation.

“Connectivity, ontologies, and knowledge maps must be adaptable to the changing environment. Gone are the days of static tables and reports as decision foundations. AI and knowledge maps help build that adaptability. Investing time getting that foundation right is a very important step.”

Jim Jacob, Executive Director of Digitalization, Advanced Analytics, and Artificial Intelligence, Cummins

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review*, the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the US and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. This content was researched, designed, and written by human writers, editors, analysts, and illustrators. This includes the writing of surveys and collection of data for surveys. AI tools that may have been used were limited to secondary production processes that passed thorough human review.

About Neo4j

Neo4j is the knowledge layer for enterprise AI. Powered by the Neo4j Graph Intelligence Platform, it includes enterprise-ready knowledge graphs for accurate, explainable, and governed AI; the most comprehensive, trusted, and easy-to-deploy graph capabilities across any environment and data source; and an unmatched ecosystem trusted by 84 of the Fortune 100 and supported by the world's largest graph community. Intelligence that works. Results that matter.



Endnotes

1. "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," Gartner, June 25, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
2. "An AI agent deleted a company's entire database in 9 seconds – then wrote an apology," Euronews, April 28, 2026, <https://www.euronews.com/next/2026/04/28/an-ai-agent-deleted-a-companys-entire-database-in-9-seconds-then-wrote-an-apology>
3. "4 Famous AI Fails (& How to Avoid Them)," Monte Carlo, November 11, 2025, <https://montecarlo.ai/blog-famous-ai-fails>
4. "Taco Bell Pumps Brakes on AI Drive-Through After Viral Fails," The Tech Buzz, September 3, 2025, <https://www.techbuzz.ai/articles/taco-bell-pumps-brakes-on-ai-drive-through-after-viral-fails>

Illustrations assembled by Tim Huxford with elements from Shutterstock.

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance by any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2026. All rights reserved.

To cite this report, please use: "Connecting AI agents to enterprise knowledge," MIT Technology Review Insights and Neo4j, October 2026.



MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com