

Leading data-driven organizations balance protection and access as AI powers ahead.

Delivering securely on data and AI strategy



Most organizations feel the imperative to keep pace with continuing advances in AI capabilities, as highlighted in a recent **MIT Technology Review Insights report**. That clearly has security implications, particularly as organizations navigate a surge in the volume, velocity, and variety of security data. This explosion of data, coupled with fragmented toolchains, is making it increasingly difficult for security and data teams to maintain a proactive and unified security posture.

Data and AI teams must move rapidly to deliver the desired business results, but they must do so without compromising security and governance. As they deploy more intelligent and powerful AI capabilities, proactive threat detection and response against the expanded attack surface, insider threats, and supply chain vulnerabilities must remain paramount. “I’m passionate about cybersecurity not slowing us down,” says Melody Hildebrandt, chief technology officer at Fox Corporation, “but I also own cybersecurity strategy. So I’m also passionate about us not introducing security vulnerabilities.”

That’s getting more challenging, says Nithin Ramachandran, who is global vice president for data and AI at industrial and consumer products manufacturer 3M. “Our experience with generative AI has shown that we need to be looking at security differently than before,” he says. “With every tool we deploy, we look not just at

Key takeaways

- 1 Data high achievers – surveyed organizations that are successfully delivering on their data strategy – put security first. Among multiple data challenges, this group puts their primary focus on three: data lineage, user authentication, and security complexity.
- 2 The high achievers use generative AI more actively than others to bolster security. Around half (48%) of this group, compared with 37% of the overall survey sample, have integrated generative AI into threat and fraud detection processes, for example.
- 3 Technology CXOs deem establishing clear governance and security a top imperative as they deploy agentic AI. Among the challenges is ensuring that authentication provides agents with controlled but sufficient access to the data they need to take action.

its functionality but also its security posture. The latter is now what we lead with.”

Our survey of 800 technology executives (including 100 chief information security officers), conducted in June 2025, shows that many organizations struggle to strike this balance. Those that do best in delivering on data strategy – the 12% of the sample that we term data high achievers – say security and governance-related

“Our experience with generative AI has shown that we need to be looking at security differently than before. With every tool we deploy, we look not just at its functionality but also its security posture. The latter is now what we lead with.”

Nithin Ramachandran, Global Vice President for Data and AI, 3M

factors are the key challenges they face in doing so. Of this select group, 46% cite difficulties tracing data lineage as a major challenge in delivering on data strategy, 42% cite challenges with user authentication, and 40% point to security complexity (see Figure 1). The survey sample as a whole, by contrast, cites shortages of skilled data engineers and scientists more than other hindrances to delivering on data strategy.

It is a similar case when it comes to scaling generative AI. While the survey group as a whole says a skills gap is their chief pain point, for nearly half of the high achievers (48%), security and IP risks are bigger ones (see Figure 2).

While difficult to draw a firm conclusion from these findings, they certainly suggest that the surveyed organizations see a connection between strong security and the successful implementation of data and AI strategy. Overall, 29% of respondents view robust security as a primary data and AI success factor.


Agentic AI, security, and governance

As AI capabilities improve, the security and governance challenges of using them are evolving. AI code-generation tools, for example, have made the lives of data engineers and application developers easier, but they can heighten the risk of insecure or poor-quality code being replicated at scale. “We’ve had success with AI-generated code, and the velocity is increasing, so we can’t have a bottleneck at infrastructure or at

“Our approach to deploying agentic AI is intentional. We focus on the intersection of high-feasibility and high-impact use cases, always leading with a responsible AI mindset.”


Rani Jonhson, Chief Information Officer,
Workday

Figure 1: Data high achievers cite security and governance-related factors as the main challenges they face in delivering on data strategy



Source: MIT Technology Review Insights survey, 2025

Figure 2: Security and IP risks are the biggest pain point for data high achievers when seeking to scale generative AI use cases (Top responses)



Source: MIT Technology Review Insights survey, 2025

deployment,” says Hildebrandt. “We need proactive guardrails to get AI code into production. This challenge, and opportunity, will become even more acute as agentic AI gains ground and we increasingly have machines writing code.”

Organizations are indeed starting to contend with security-related and other challenges associated with agentic AI – systems capable of autonomous decision-making and action. Just under one-fifth (19%) of respondents say their organizations have begun using agentic AI models (see Figure 3), and 68% say they will start doing so within the next two years.

Ensuring clear governance and security looms large as a challenge for organizations in supporting their agentic AI models. This ranks second of around a dozen potential agentic AI difficulties, cited by 38% (see Figure 4).


In many respects, the governance and security considerations of agentic AI are not materially different from those of generative AI. According to Jonny LeRoy, senior vice president and CTO at equipment producer W.W. Grainger: “Like other forms of AI, agentic AI will work really well if you have good playbooks, good knowledge management, and good policies.”

It is the potential dangers of agents autonomously making decisions that explain why several executives we interviewed say their organizations are currently experimenting with agentic AI in low-risk use cases. This is the case at W.W. Grainger, says LeRoy: “Our path at the moment is very much human in the loop. We start with an agent making suggestions and a human saying

“Agentic AI demands well-structured, high-context, explainable features, which means lineage, governance, and semantic clarity become even more critical.”

Sejung Lee, Chief Data Officer, KT

Figure 3: 1 in 5 organizations are deploying agentic AI to some extent, while two-thirds are deploying generative AI



Source: MIT Technology Review Insights survey, 2025


Figure 4: Governance and security is the second highest difficulty organizations currently encounter, or are likely to encounter, in supporting agentic AI



Source: MIT Technology Review Insights survey, 2025

“Like other forms of AI, agentic AI will work really well if you have good playbooks, good knowledge management, and good policies.”

Jonny LeRoy, Senior Vice President and Chief Technology Officer,
W.W. Grainger



yes or no. As we gain confidence, that will start changing, but with a new set of capabilities you’ve got to build your way up to it.”

Similarly, Rani Johnson, chief information officer at enterprise software provider Workday, says her team is currently prioritizing high-feasibility, low-data, and low-risk use cases. “Our approach to deploying agentic AI is intentional. We focus on the intersection of high-feasibility and high-impact use cases, always leading with a responsible AI mindset,” she says. “The difference in our governance threshold between an agent recommending an IT fix and one handling sensitive HR or benefits data is immense. For high-stakes scenarios, we invest the time in building auditable controls and governance frameworks to protect our people.”

The aforementioned survey results also indicate that agentic AI has implications for data governance (see Figure 4). “Agentic AI demands well-structured, high-context, explainable features – which means lineage, governance, and semantic clarity become even more critical,” says Sejung Lee, chief data officer at South Korean telecom provider KT. “We’re investing in responsible AI design, feature engineering pipelines, and ongoing validation frameworks to keep agents grounded and auditable. But we know these systems are dynamic, and governance must evolve with them.”

Authenticating AI at speed


As mentioned earlier, 42% of high achievers cite issues with user authentication as a major challenge to delivering on their data strategy (see Figure 1). If authentication processes are overly complex – for example, requiring multiple approvals or the use of

Generative AI for tighter security

The data high achievers set themselves apart from the rest of the surveyed organizations in another respect. That is in their use of generative AI to bolster cybersecurity. When asked how they are integrating generative AI capabilities into operations and workflows, 48% of high achievers, compared with 37% of the overall sample, cite security uses such as threat and fraud detection (see Figure 5).

Generative AI is, of course, a double-edged sword when it comes to security. Attackers use it to create malware, find vulnerabilities in code, and create convincing content for phishing emails, for example. But organizations are also using generative AI to good effect in several ways. These include analysis of threat vectors and identification of emerging threats, simulation of attacks for training purposes, and automation of routine security processes.

Figure 5: Nearly half of data high achievers have integrated generative AI into security (Top responses)



Source: MIT Technology Review Insights survey, 2025

non-standard credentials – or poorly integrated with other identification systems, they can hinder data access and, ultimately, the scaling of AI use cases. Well-designed authentication, by contrast, can enable scaling by providing for reliable and secure data access.

Authentication is on the mind of executives as they deploy agentic AI more widely. One challenge they are deliberating on is how to classify AI agents. Some wonder whether agents can be provided access to systems in the same way that humans are. “That’s our starting point, to model agents like they are humans in our enterprise,” says Hildebrandt at Fox Corporation. “We have principles like least privilege for how we provision access to employees. In the same way, we shouldn’t be granting root admin to AI agents in our systems. We wouldn’t loosely grant that kind of access to a human.”

Agent authentication is necessary to ensure they have the data needed to take appropriate actions, says Johnson. “The key to unlocking the true power of agentic AI is not just functionality but safely federating access to the data it needs,” she says. “An agent is limited if it only operates within a single system. Our focus must be on architecting a unified data layer and granular security model so

agents can safely access and act on data across multiple systems, which is what gives them real enterprise value.”

More threats, more protection


CXOs can be certain that AI will continue to advance, bringing new security challenges with it. Like AI models themselves, attacks will be increasingly multi-modal. Malware will adapt to new defenses dynamically and rapidly. Attackers will use AI’s capabilities to more efficiently identify vulnerabilities and create more convincing content to gain illicit access to models and systems.

But AI advances will also create new opportunities to address these and other threats. Security teams will use AI to improve threat detection and prevention, automate incident response, and expand authentication options, for example.

Arguably the biggest advantage that AI will give security teams is the ability to shift from purely reactive mode to a more proactive and predictive form of security. This will improve organizations’ chances of finally delivering the business results from AI that have eluded them.

“We need proactive guardrails to get AI code into production. This challenge, and opportunity, will become even more acute as agentic AI gains ground and we increasingly have machines writing code.”

Melody Hildebrandt, Chief Technology Officer, Fox Corporation



“Delivering securely on data and AI strategy” is an executive briefing paper by MIT Technology Review Insights. Denis McCauley was the author of the report, Virginia Wilson was the editor, and Nicola Crepaldi was the publisher. MIT Technology Review Insights has independently collected and reported on all findings contained in this paper.

We would like to thank the sponsor, Databricks, as well as the following experts for their time and insights:

Melody Hildebrandt, chief technology officer, Fox Corporation

Rani Johnson, chief information officer, Workday

Sejung Lee, chief data officer, KT

Jonny LeRoy, chief technology officer, W.W. Grainger

Nithin Ramachandran, global vice president, data and AI, 3M

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.

From the sponsor

Databricks is the Data and AI company. More than 20,000 organizations worldwide — including Block, Comcast, Condé Nast, Rivian, Shell, and over 60% of the Fortune 500 — rely on the Databricks Data Intelligence Platform to take control of their data and put it to work with AI. Databricks is headquartered in San Francisco, with offices around the globe, and was founded by the original creators of Lakehouse, Apache Spark™, Delta Lake, MLflow, and Unity Catalog. To learn more, follow Databricks on [X](#), [LinkedIn](#), and [Facebook](#).



Illustrations

Cover art from Adobe Stock and spot illustrations created from Adobe Stock and The Noun Project.

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 on any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2025. All rights reserved.



MIT Technology Review Insights

www.technologyreview.com
insights@technologyreview.com