MIT Technology Review Insights

Produced in partnership with

databricks

What it takes to deliver on data and AI strategy.

Building a high-performance data and AI organization

2nd edition

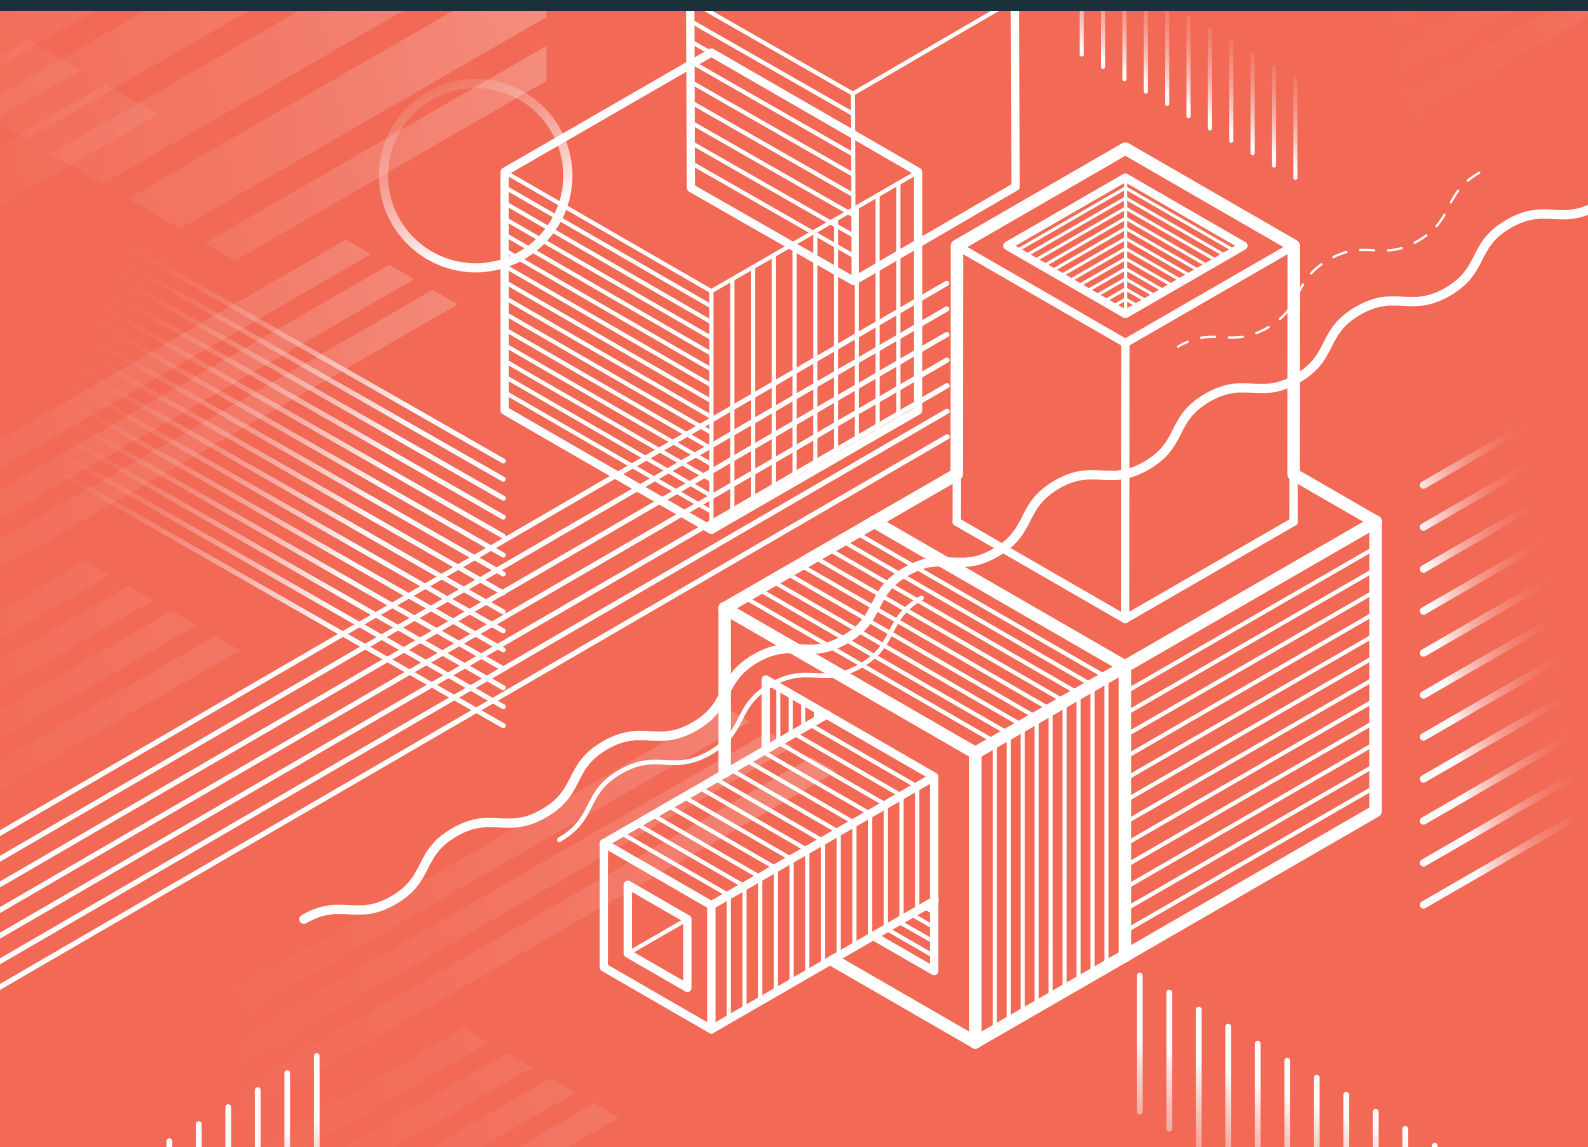

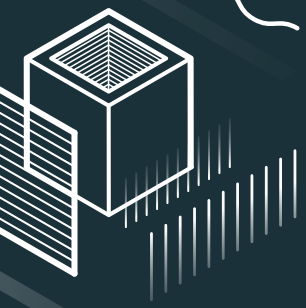

Preface

“Building a high-performance data and AI organization, 2nd edition” is an MIT Technology Review Insights report sponsored by Databricks. This report, based on survey research and executive interviews, seeks to understand how data management technologies and practices for AI have advanced in organizations since the first edition of this report in 2021. 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:

Christopher d’Arcy, Chief Data & AI Officer and Managing Director, E.ON Digital Technology

Rafael Cavalcanti, Chief Data and AI Officer, Bradesco

Dan Chaddock, Chief Technology Officer, Foundational Technology and Platform, Navy Federal Credit Union

Helius Guimaraes, Chief Data & AI Officer, Fonterra Co-operative Group

Melody Hildebrandt, Chief Technology Officer, Fox Corporation

Rani Johnson, Chief Information Officer, Workday

Irfan Khan, President and Chief Product Officer, SAP Data and Analytics

Sejung Lee, Chief Data Officer, Korea Telecom

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

Bastien Parizot, Senior Vice President, Technology and Digital, Reckitt

Shiyi Pickrell, Senior Vice President for Data and AI, Expedia Group

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

Bharathi Viswanathan, Chief Digital and Information Officer, Suntory Beverages International

Murali Vridhachalam, Vice President and IT Head of Cloud, Data, and AI, Gilead Sciences

Methodology

In June 2025, MIT Technology Review Insights, in collaboration with Databricks, conducted a global survey of 800 chief information officers (CIOs), chief technology officers (CTOs), chief data and analytics officers (CDAOs), and other senior data and technology executives. The survey respondents work in organizations that are headquartered in 12 countries. There are eight sectors represented in the sample, as well as the category of digital native companies. All respondents work in organizations earning $500 million or more in annual revenue.

In addition to the quantitative research from the survey, a series of in-depth interviews with CIOs, CTOs, chief data officers (CDOs), and chief artificial intelligence officers (CAIOs) of large private and public sector organizations offer firsthand insights into data management technologies and practices for AI in organizations.

CONTENTS

01 Executive summary 3

02 Racing to keep pace 5

Adaptability struggles 5

Databricks perspective: Unifying data, analytics, and AI for business impact 7

The priorities ahead 10

FOX's Sports AI chatbot 11

03 Rise of the agents 12

Generative AI scaling challenges 13

Enter agentic AI 14

Challenges of agentic AI 15

Testing the waters with agentic AI 17

04 AI-powered data management 18

Data intelligence 19

Avoiding agent chaos 20

05 Freedom to act 21

Cloud flexibility 22

06 Conclusion 23

01 Executive summary

Four years is a lifetime when it comes to artificial intelligence. Since the first edition of this study was published in 2021, AI's capabilities have been advancing at speed, and the advances have not slowed since generative AI's breakthrough. For example, multimodality – the ability to process information not only as text but also as audio, video, and other unstructured formats – is becoming a common feature of AI models. AI's capacity to reason and act autonomously has also grown, and organizations are now starting to work with AI agents that can do just that.

Amid all the change, there remains a constant: the quality of an AI model's outputs is only ever as good as the data that feeds it. Data management technologies and practices have also been advancing, but the second edition of this study suggests that most organizations are not leveraging those fast enough to keep up with AI's development. As a result of that and other hindrances, relatively few organizations are delivering the desired business results from their AI strategy. No more than 2% of senior executives we surveyed rate their organizations highly in terms of delivering results from AI.

To determine the extent to which organizational data performance has improved as generative AI and other AI advances have taken hold, MIT Technology Review Insights surveyed 800 senior data and technology executives. We also conducted in-depth interviews with 15 technology and business leaders.

Key findings from the report include the following:

- **Few data teams are keeping pace with AI.** Organizations are doing no better today at delivering on data strategy than in pre-generative AI days. Among those surveyed in 2025, 12% are self-assessed data "high achievers" compared with 13% in 2021. Shortages of skilled talent remain a constraint, but teams also struggle with accessing fresh data, tracing lineage, and dealing with security complexity – important requirements for AI success.

- **Partly as a result, AI is not fully firing yet.** There are even fewer "high achievers" when it comes to AI. Just 2% of respondents rate their organizations' AI performance highly today in terms of delivering measurable business results. In fact, most are still struggling to scale generative AI. While two thirds have deployed it, only 7% have done so widely.

- **Data strategies are not highly adaptable.** Most surveyed executives lack confidence in the ability of their

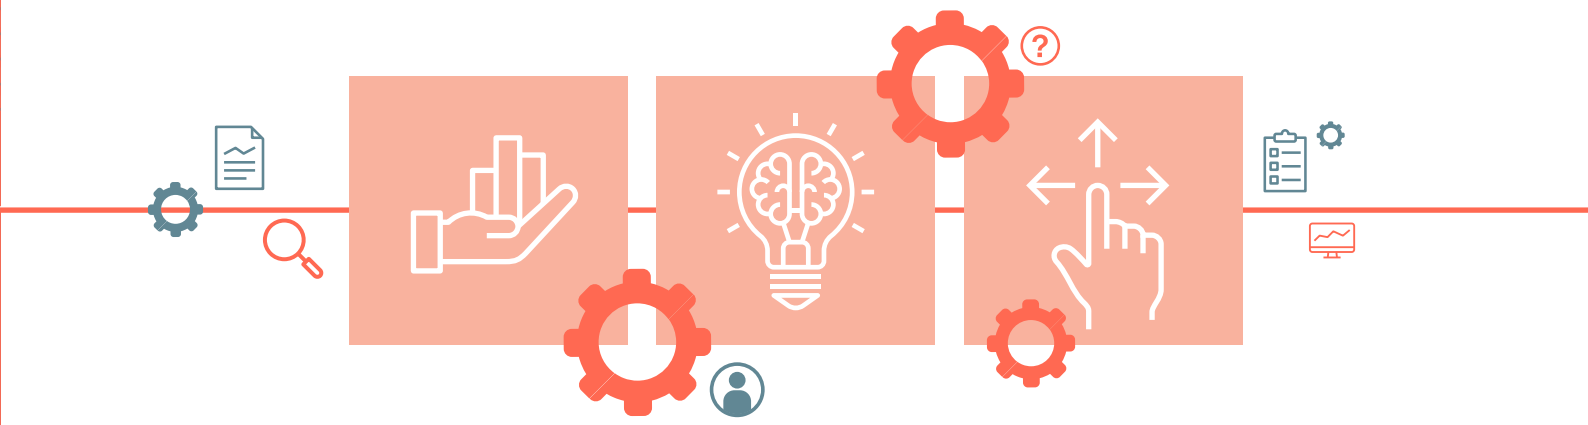

The vast majority of enterprises are deploying AI tools and data intelligence. Now, agentic AI is raising the stakes again. Organizations are taking their time to ensure robust governance and secure ample high-quality, explainable data to support agentic AI's advanced capabilities before building momentum.

organizations' data strategies to adapt to future AI advances. Less than half (46%) say data strategy is largely or highly capable of doing so. Difficulties aligning data and AI strategies are likely a significant factor, manifestations of which include separate data and AI governance models and a lack of unified technology platforms.

- **Agentic AI raises the stakes for data teams.** Machines taking actions puts an extra onus on data and AI leaders to ensure clear and robust data governance and security. Organizations are moving slowly so far with agentic AI and are working to secure a sufficiency of high-quality and explainable data to support its advanced capabilities.

- **AI tools foreshadow change to data management itself.** Adoption of AI-powered tools is widespread among the surveyed organizations: 67% use such tools now, and the rest will be adopting them in the next one to two years. Data team efficiency will benefit from automation of tasks, but sourcing will require discipline to avoid added toolset complexity.

- **Data intelligence will help data teams better understand their assets.** A majority of respondents (69%) say their organizations make use of data intelligence today, with most of the rest saying they will start doing so soon. Respondents say they will use it to improve data governance, connect data silos, improve their organization of data, and analyze more of it.

02 Racing to keep pace

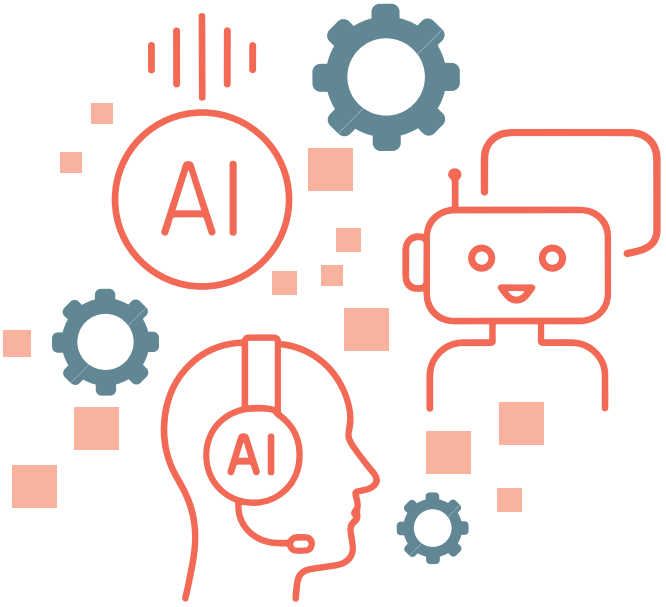

The relentless progress of AI has placed enormous pressure on enterprise data teams to improve their performance. The demand is to deliver the high-quality, well-curated, and trustworthy data that today's AI models require – and to do so, increasingly, in real time.

Judging by the survey results, data performance in organizations leaves much to be desired. We asked the respondents to rate data performance in terms of their organization's success in delivering on data strategy with measurable business results. Against that criterion, just 12% provide a high rating (nine or 10 on a one to 10 scale). We term these as data "high achievers".

The percentage of "high achievers" remains essentially unchanged from when we asked this question in our pre-generative-AI study in 2021; the figure then was 13%. And "low achievers" (those rating their company's performance at six or below) are more numerous today than pre-generative AI (18% versus 12%) (see Figure 1a). Taken together, these figures suggest that most organizations' data teams are struggling to deliver what their AI models and use cases require.

That conclusion is supported by responses to the comparable survey question about delivering on AI strategy. Today, "high achievers" in AI performance amount to just 2% of those surveyed, and "low achievers" to 71% (see Figure 2). It is not a stretch, then, to suggest that organizations' data struggles are impeding them from generating valuable business results with their AI use cases.

The relentless progress of AI places enormous pressure on enterprise data teams to deliver the well-curated and trustworthy data today's AI models require – and to do so, increasingly, in real time.

Adaptability struggles

Why are so many organizations struggling with data and AI? One reason is that many cannot adapt their data strategy and infrastructure rapidly enough to keep pace with AI advances. "AI, especially generative AI, is moving extremely fast, but data strategy doesn't move at the same pace," says Sejung Lee, CDO at South Korean telecom provider KT. "For example, we at KT still operate with over 1,800 legacy systems, and many of them were never designed with AI in mind."

Indeed, survey respondents give a lukewarm assessment of the adaptability of their data strategy: 46% say it is "largely" or "highly capable" of adapting to future

Figure 1a: The extent to which organizations are successfully delivering on their enterprise data strategy remains broadly unchanged, pre-generative AI and today
(Self-assessed rating on a 1 - 10 scale where 10 = succeeding)

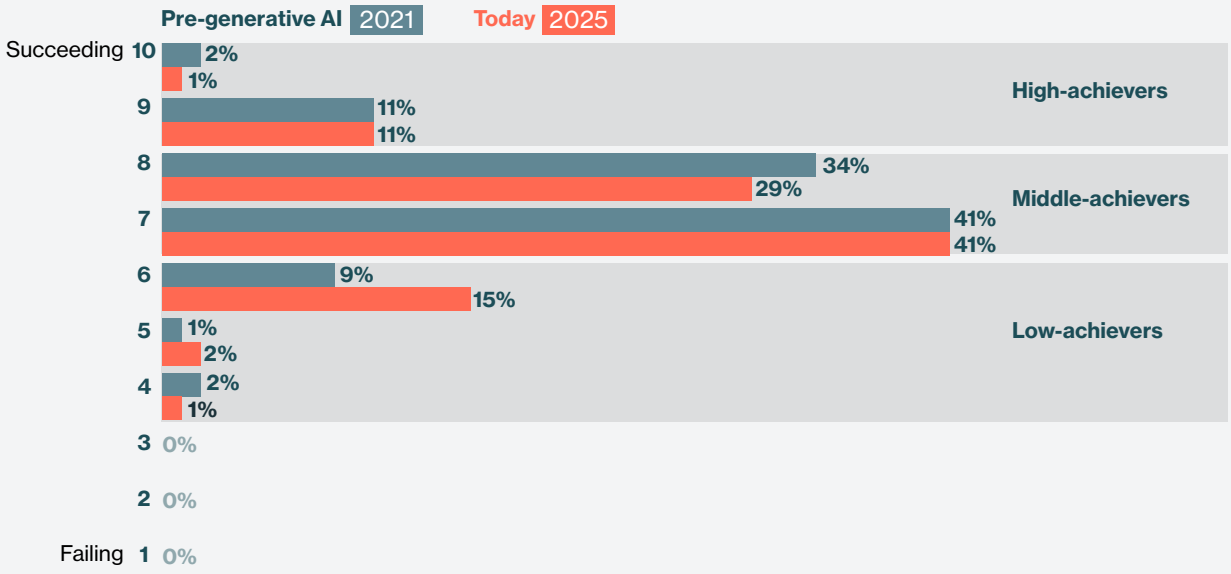

Figure 1b: Data "high achievers" in 2025, by organization size (annual revenue) and industry
(% of respondents)

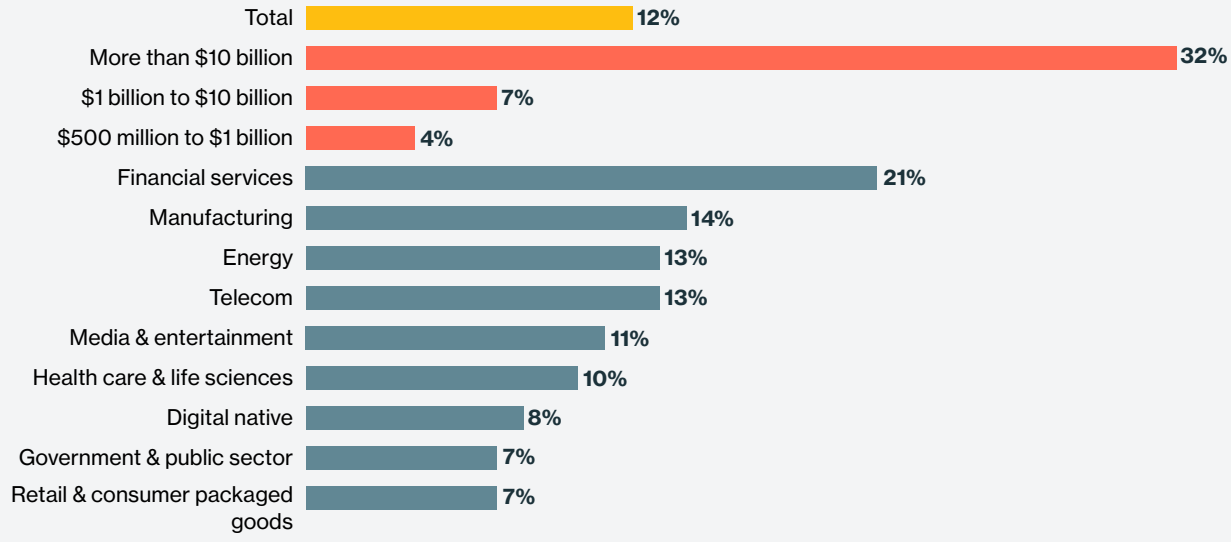

Source: MIT Technology Review Insights survey, 2025

As 65% of organizations deploy generative AI, the largest come out on top for data performance – 32% of $10 billion+ companies were ranked high achievers.

Partner perspective

Databricks perspective:

Unifying data, analytics, and AI for business impact

The survey findings reinforce a critical truth in today's AI-driven economy: organizations that achieve the greatest business impact from AI are those that unify their data, analytics, and AI on a single, open, and collaborative platform. Such an approach enables higher data quality, streamlined governance, and faster deployment of AI models—capabilities that grow in importance as generative and agentic AI mature. Databricks's AI positioning, particularly through Agent Bricks, centers on addressing three key customer needs:

- First, accuracy is context-dependent, meaning customers require agents that deeply understand their business and specific requirements from their unique enterprise data. This necessitates an AI solution that not only understands the data but can measure the accuracy of the AI results from the data and continuously improves quality.
- Second, recognizing that AI moves quickly, customers need the flexibility to use models from any provider, make necessary changes, and serve them at the lowest possible cost to support a broad range of use cases, ensuring they are not left behind by rapid innovation.
- Third, it is crucial to bring the model to the data to ensure that governance and security remain intact. Customers are keen to avoid new AI-related risks such as shipping data outside their boundaries or becoming locked into closed technologies.

This focus on quality and openness, coupled with strong governance, accelerates the path from experimentation to scaled production. As AI technologies evolve, the competitive advantage will belong to organizations that align data and AI strategies to measurable outcomes, not just technology adoption. Platforms that connect teams, reduce fragmentation, and provide trusted, explainable AI at scale will be the foundation for delivering sustained business value. Databricks, with its open and unified Data Intelligence Platform, is uniquely positioned to help enterprises achieve this balance—turning data and AI investments into measurable impact.

advances in AI. The majority (42%) say it is at best "somewhat capable" of adapting (see Figure 3). Not surprisingly, data "high achievers" have greater confidence on this score: 80% deem their data strategy as largely or highly capable of adapting.

Another reason organizations struggle is likely a lack of alignment between data and AI strategies. "We learned very quickly that those strategies have to be closely intertwined," says Rani Johnson, CIO at enterprise software provider Workday. "When we started with generative AI, if we got bad results, we found it was mainly due to bad data." The two strategies are similarly

> "We learned very quickly that data and AI strategies have to be closely intertwined."
>
> Rani Johnson, Chief Information Officer, Workday

intertwined at equipment producer W.W. Grainger, according to Jonny LeRoy, its CTO. "We're doing a lot of work at the team level to bring machine learning, engineering, architecture, product, and design together so that they're working in the same ways."

Figure 2: AI strategy delivery is advancing gradually, with few organizations reporting full success
(Self-assessed rating on a 1 - 10 scale where 10 = succeeding)

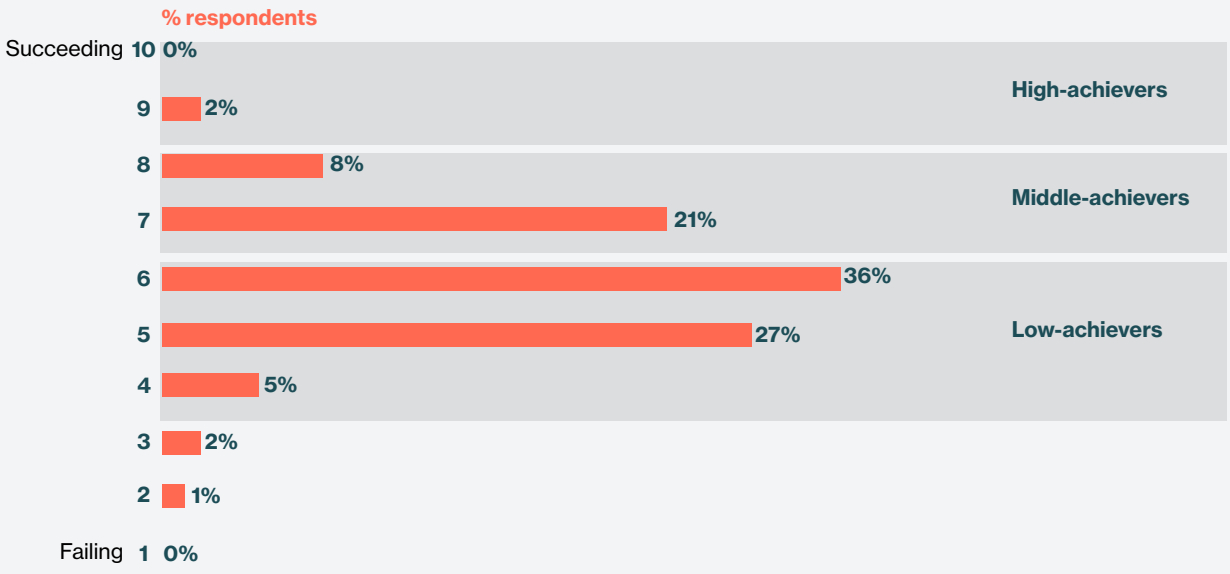

Source: MIT Technology Review Insights survey, 2025

Figure 3: 46% of respondents say their organization's enterprise data strategy is largely or highly capable of adapting to future advances in AI
(% of respondents)

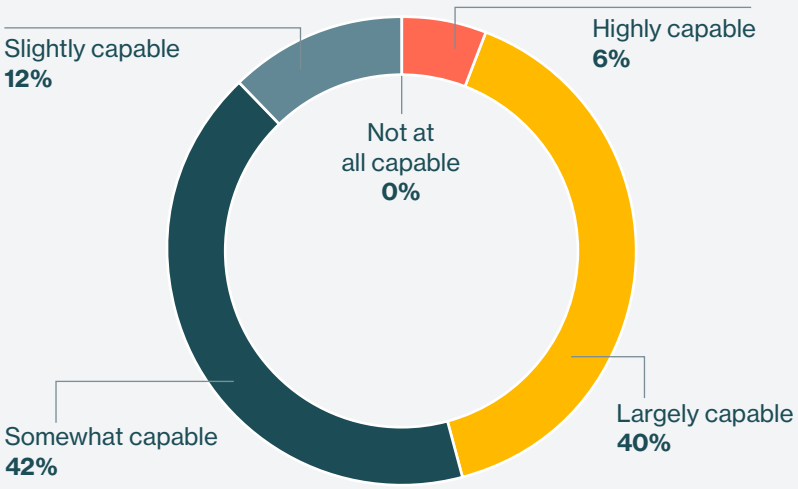

Source: MIT Technology Review Insights survey, 2025

"Generative AI made clear that organizations that combine their data and AI strategies can more easily deploy sophisticated AI approaches than those where data and AI are managed separately."

Rafael Cavalcanti, Chief Data and AI Officer, Bradesco

Figure 4: Talent shortages, access to trackable, fresh data, and security complexity are primary challenges to delivering on data strategy
(% of respondents)

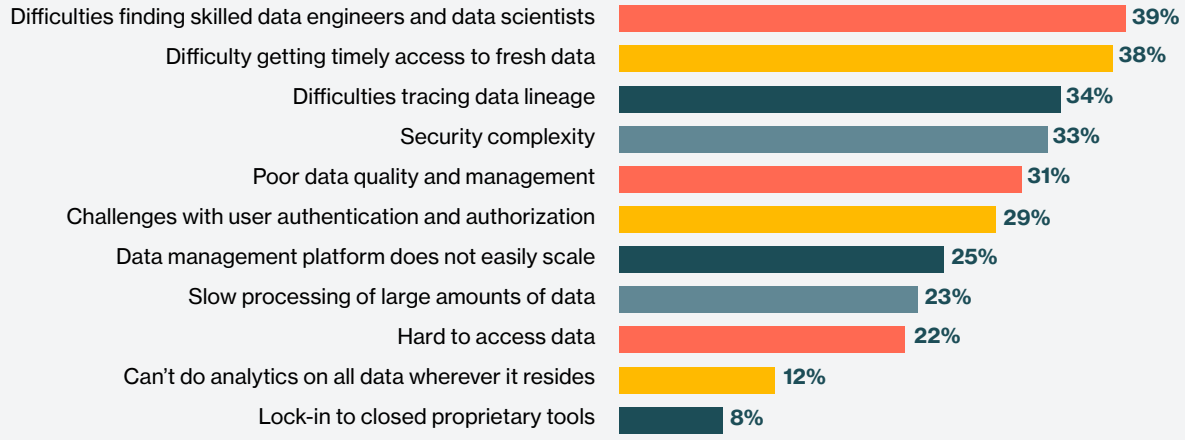

Source: MIT Technology Review Insights survey, 2025

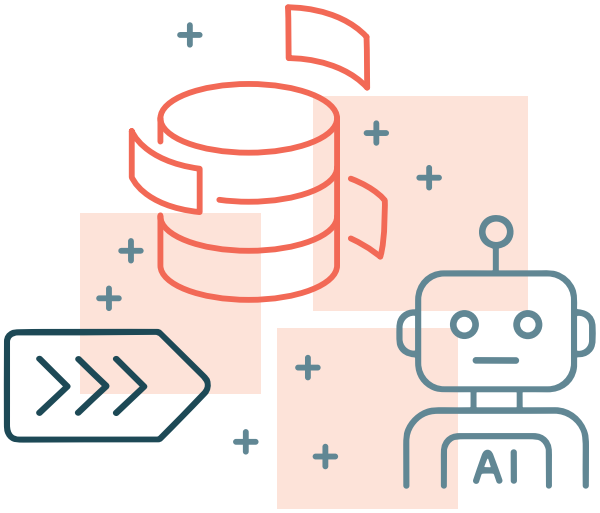

"With all the investment you put into your infrastructure layers, robust governance really pays off when new technologies like generative AI come along."

Shiyi Pickrell, Senior Vice President for Data and AI, Expedia Group

When generative AI began to make waves in late 2022, Bradesco, a Brazil-based bank, undertook a refresh of its data strategy, says Rafael Cavalcanti, the company's CDAIO. One result of that was a decision to combine Bradesco's data assets with its AI initiatives, and its data and AI strategies have since been intertwined. "Generative AI made clear that organizations that do this can more easily deploy sophisticated AI approaches than those where data and AI are managed separately," says Cavalcanti.

For the highest AI achievers in the survey, the factor cited most frequently in their success is "strong business alignment of AI and data strategies tied to measurable business outcomes." And for the highest data achievers, that alignment ranks second in importance only to improved data quality.

Three weaknesses do more than other factors to hinder alignment of data and AI strategy. They are: separate data and AI governance models (cited by 41%), the lack of unified technology platforms (37%), and different approaches taken by data and AI teams to measuring return on investment (cited by 32%).

When it comes to implementing data strategy, shortages of skilled data scientists and engineers continue to hamper organizations. In the survey group as a whole, this

Figure 5: The most important enterprise-wide data and AI strategy initiatives for survey respondents over the next two years
(Top responses; % of respondents)

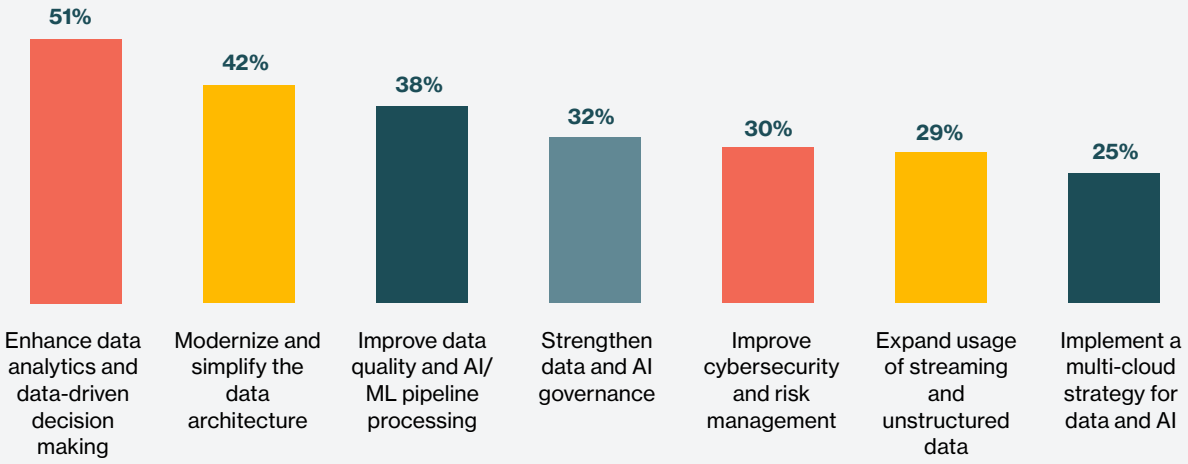

Source: MIT Technology Review Insights survey, 2025

tops the respondents' list of challenges to delivering on data strategy (cited by 39%). The other big challenges are difficulty getting timely access to fresh data (38%), difficulties tracing data lineage (34%), and security complexity (33%) (see Figure 4).

Accessing fresh data was a time-consuming process at Gilead Sciences, says Murali Vridhachalam, its head of cloud, data, and AI in the IT organization. "We used to buy more than 250 large sets of pharma data from more than 100 vendors, and we used to have to manually scan those for malware, something that took several days. Thanks to our investment in cloud automation technologies, one of the largest data sets, a 40-terabyte file, now takes just two to three hours to scan and load the data and be ready for use." Automation has also helped to streamline lineage tracking, says Vridhachalam.

The priorities ahead

Talent shortages aside, the data and AI priorities of surveyed organizations over the next two years provide an indication of how they aim to improve performance in both areas. At the top of their list is enhancing analytics and data-driven decision making (cited by 51%). Next are modernizing and simplifying the data architecture (42%), followed by improving data quality and AI/machine learning (ML) pipeline processing (38%). Other priorities include strengthening data and AI governance, improving cybersecurity and risk management, and expanding the use of streaming and unstructured data (see Figure 5).

> "Thanks to our investment in cloud automation technologies, it takes just two to three hours to scan [a 40-terabyte file of pharma data] and be ready for use."
>
> Murali Vridhachalam, Vice President and IT Head of Cloud, Data, and AI, Gilead Sciences

Despite the relatively long history of data analytics in enterprise use, many organizations have only scratched the surface of the ability of these tools to deliver insights. According to Bharathi Viswanathan, chief digital and information officer at Suntory Beverages, companies like hers have tended to focus on what analytics can do for the front end of the business. "While we have made great progress on commercial and topline-focused analytics,

FOX's Sports AI chatbot

FOX Sports has long been a leader in live sports broadcasting, and its in-depth coverage of the game, both on and off the field, has supplied fans with stats, highlights, color commentary, and player perspectives. With the launch of Sports AI, FOX Sports has recently gained the powers of an AI chatbot, able to answer users' questions about all the sports that FOX covers.

According to Melody Hildebrandt, CTO at Fox Corporation, the idea behind Sports AI is to leverage the variety of reporting and other content that FOX journalists and other experts deliver across a range of sports. "Users can ask Sports AI, for example, who's going to win the NFL championship this year," Hildebrandt explains. "The app responds with a prediction and summary that's based on the text written by our journalists, the on-air commentary of our experts, and other coverage of the sport. It's delivered through a multi-modal AI model."

Among the biggest challenges to making Sports AI work as intended was transforming how the app searches for content. "We had to create a data foundation that allowed us to essentially vectorize our content," says Hildebrandt. "We needed to replace the basic keyword search with semantic search." So, her team swapped out the entire backend of the app.

Now, says Hildebrandt, FOX Sports has a data foundation that helps the app understand all its content. "The data pipelines are now all coming into a system that allows us to invoke different models for different use cases without having to start from scratch each time. That foundation has also led to a much greater speed of data delivery."

"The data pipelines are now all coming into a system that allows us to invoke different models for different use cases without having to start from scratch each time. That foundation has also led to a much greater speed of data delivery."

Melody Hildebrandt, Chief Technology Officer, Fox Corporation

we have barely touched its potential in supply chain and procurement, HR, and many other areas of operation. There's so much more to do with analytics," says Viswanathan.

Simplifying the data architecture and centralizing governance models, meanwhile, are big priorities for online travel agency Expedia Group, according to Shiyi Pickrell, its senior vice president for data and AI. "One thing I've learned in my time here is that with all the investment you put into your infrastructure layers, robust governance really pays off when new technologies like generative AI come along," says Pickrell.

Dan Chaddock, CTO, foundational technology and platform at the Navy Federal Credit Union in the US, links decision making and transparency directly to data quality. "Data quality is the top priority for us. If we really want to understand how we're making decisions and have transparency, it all starts with that."

"Data quality is the top priority for us. If we really want to understand how we're making decisions and have transparency, it all starts with that."

Dan Chaddock, Chief Technology Officer, Foundational Technology and Platform, Navy Federal Credit Union

03 Rise of the agents

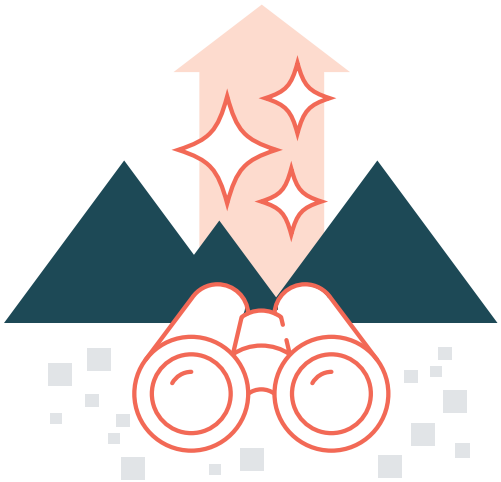

The vast media attention given to generative AI in recent years creates an impression of almost universal enterprise adoption. Our survey does not bear this out. Traditional AI – which relies on pre-programmed rules and algorithms to perform specific tasks but lacks the ability to adapt to evolving situations or generate new ideas – is pervasive. Some 97% of the surveyed respondents say their organization is using it, including 71% where it is deployed widely.

Fewer (65%) say they have deployed generative AI and the majority (58% overall) have done so only in limited fashion. Just 7% have deployed it widely (see Figure 6). The limited deployment of generative AI is yet another reason why AI strategy is not yet delivering a measurable business impact for most organizations.

The surveyed organizations have sought thus far to integrate generative AI mainly into internal operations. The top use cases cited by respondents involve internal

Figure 6: Of the 65% of organizations that have deployed generative AI, the majority have done so only in a limited way
(% of respondent organizations with limited or widescale deployment)

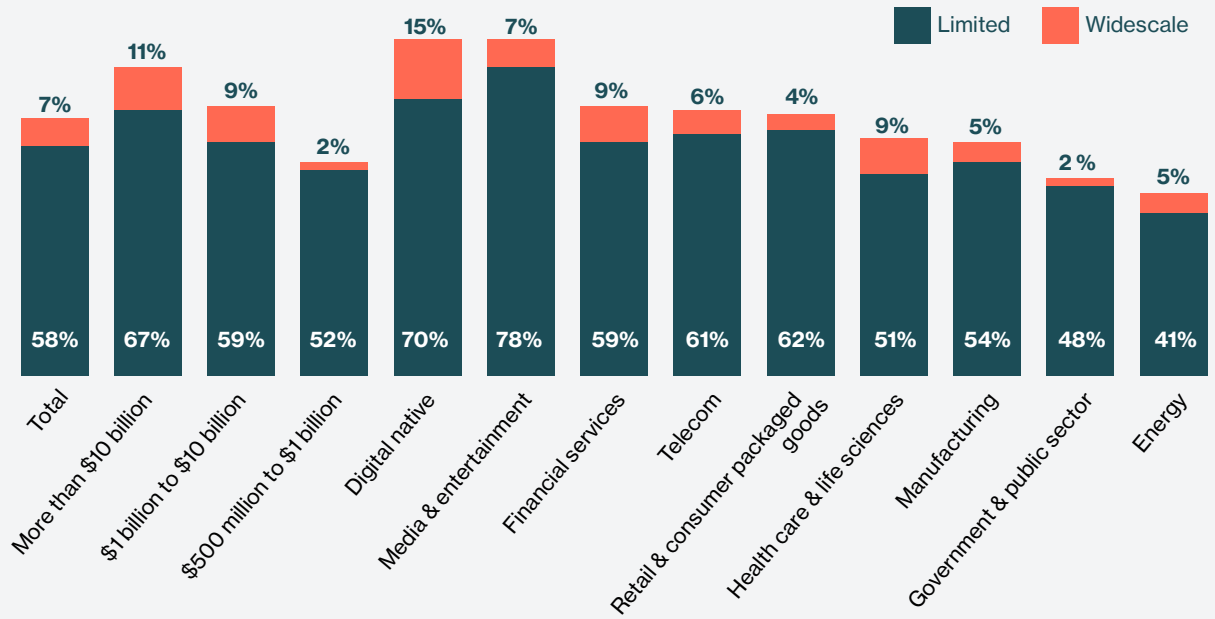

Source: MIT Technology Review Insights survey, 2025

productivity (such as summarizing documents) (cited by 52%), data analysis or decision support (47%), and compliance (41%). Fewer organizations have integrated it into external-facing operations such as customer service (38%) (see Figure 7).

Generative AI scaling challenges

The chief barrier cited by respondents to scaling generative AI is similar to that of delivering on data strategy – a skills gap (mentioned by 45%) (see Figure 8). Considerably fewer of the "high achievers" (31%), however, struggle with this. The technology leaders we interviewed acknowledge the difficulties in sourcing talent but say their organizations have found ways of adapting. "Finding data scientists and engineers is a challenge in today's market," says Bastien Parizot, senior vice president, technology and digital at Reckitt, a consumer goods company. "But we've built strong relationships with external partners, which allows us to focus our internal resources on our chosen priorities."

Figure 7: Generative AI is currently more commonly used in internal operations or workflows
(% of respondents)

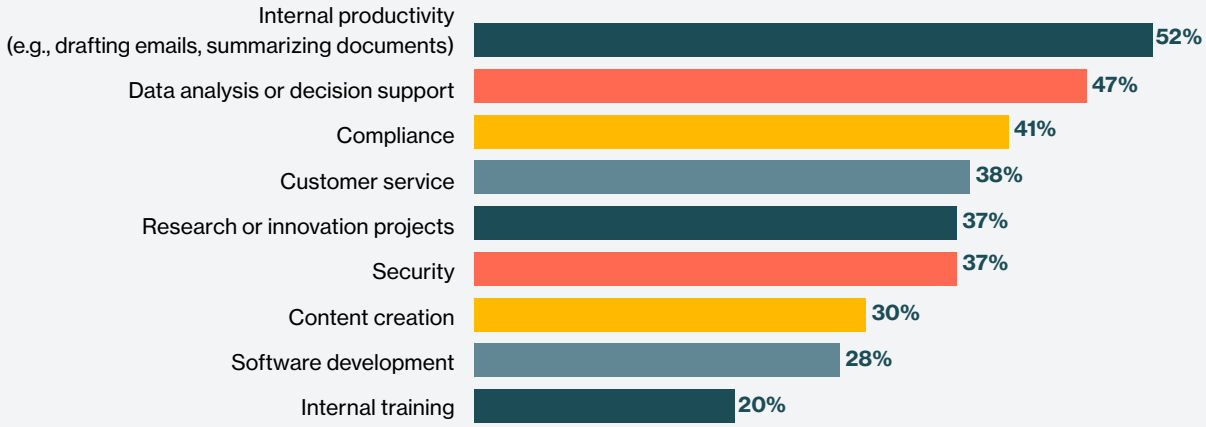

Source: MIT Technology Review Insights survey, 2025

Figure 8: Chief pain points for organizations seeking to scale generative AI use cases include a skills gap, fragmented infrastructure, and regulatory uncertainty
(% of respondents)

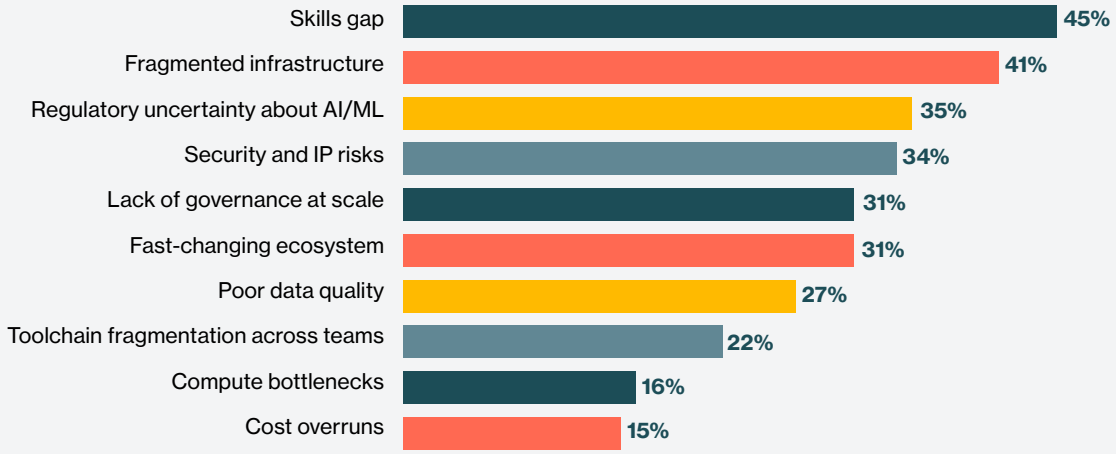

Source: MIT Technology Review Insights survey, 2025

Energy provider E.ON has also coped well with a tight talent market, says Christopher d'Arcy, CDAIO and managing director at E.ON Digital Technology. "We started building up AI skills long before the generative AI hype started. So we have a good nucleus of talented scientists and engineers who serve the whole company on AI projects," says d'Arcy.

It is a similar story at Suntory Beverages, says Viswanathan. "We have an outstanding data team that has been able to take gold AI technology and mine it," she says. "But we've invested a lot of time getting the right talent and training on our business, to spark their imagination, and basically to get them to love their jobs. When people love their jobs, you get magic."

Fragmented infrastructure is another common pain point in scaling generative AI (cited by 41% of respondents) (see Figure 8). "The architectural issues are very challenging," says Melody Hildebrandt of Fox Corporation. "The ecosystem is moving fast, and you want to respond to that, but you can't constantly be re-architecting. So it's deciding what you absolutely need and what you can live with for a period of time."

For Viswanathan, ensuring the company's core platforms are ready for AI is a focus of her team's work on

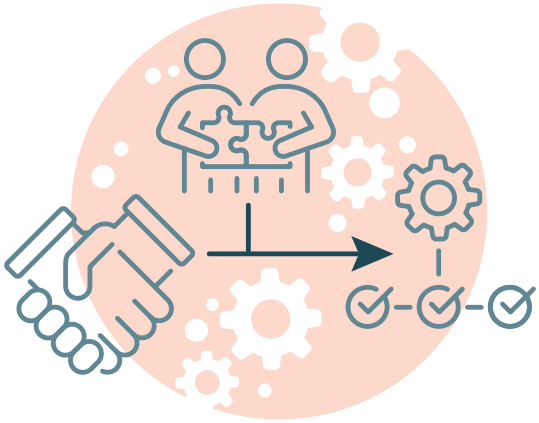

"Finding data scientists and engineers is a challenge in today's market. But we've built strong relationships with external partners, which allows us to focus our internal resources on our chosen priorities."

Bastien Parizot, Senior Vice President, Technology and Digital, Reckitt

Figure 9: 19% of organizations have deployed agentic AI, mostly to a limited extent
(% of respondent organizations with limited or widescale deployment)

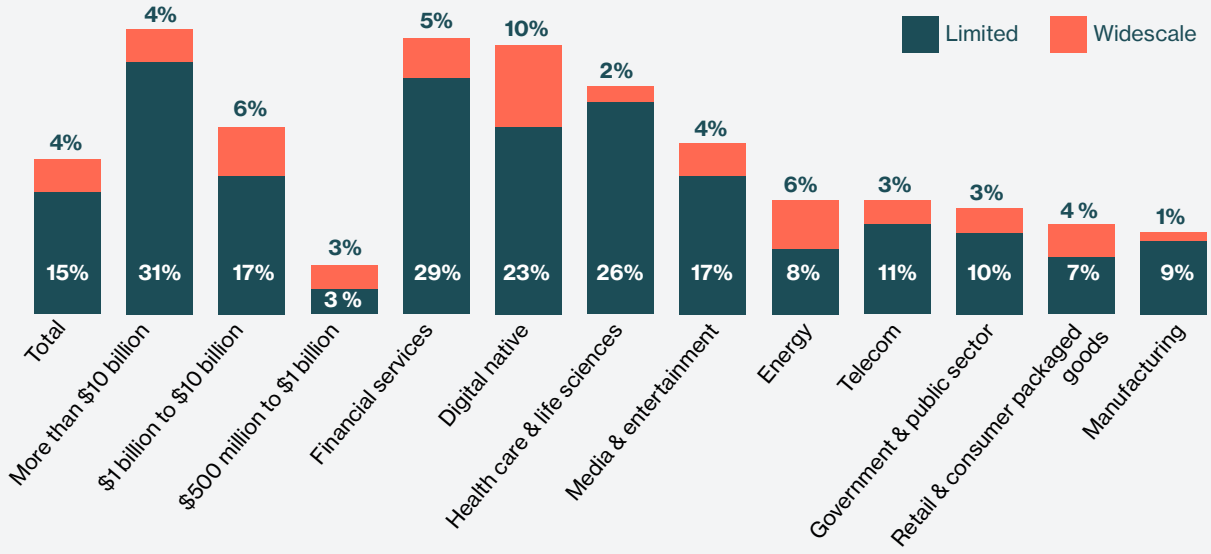

Source: MIT Technology Review Insights survey, 2025

infrastructure. “We’ve got to move the needle on both our AI platforms and our core, underlying platforms. Because AI is a major user of the latter.”

Enter agentic AI

Agentic AI – systems capable of autonomous decision making and action – ushers in the next set of challenges that data and AI teams will need to square up to. In the survey, 19% of respondents say their organizations have started using agentic AI models (see Figure 9). And 68% say their organizations will invest in developing such capabilities over the next one to two years. The data “high achievers” are well ahead of the curve on agentic AI: 46% have already begun deploying it.

As with generative AI, the surveyed organizations are principally aiming to yield gains from agentic AI in internal operations. Over half (53%) cite the opportunity to use it to enhance operational efficiency, while 45% look to improve data analysis and insight generation. Accelerating decision-making processes, automating routine or repetitive tasks, and improving IT operations (all cited by 41%) are other sought-for gains from agentic AI (see Figure 10).

“We’ll use [agentic AI] to drive efficiency across the organization in ways that deliver competitive advantage, including by improving our R&D processes,” says Nithin Ramachandran, global vice president, data and AI at industrial and consumer products manufacturer 3M. The Navy Federal Credit Union pivoted toward agentic AI around a year ago, says Chaddock, and it will eventually whittle down around 200 existing use cases. “Right now we’re using agentic AI in productivity use cases that make

> “We have an outstanding data team that has been able to take gold AI technology and mine it. But we’ve invested a lot of time getting the right talent and training on our business, to spark their imagination, and basically to get them to love their jobs. When people love their jobs, you get magic.”
>
> Bharathi Viswanathan, Chief Digital and Information Officer, Suntory Beverages International

Figure 10: When it comes to agentic AI, more respondents see opportunities in internal operations
(% of respondents)

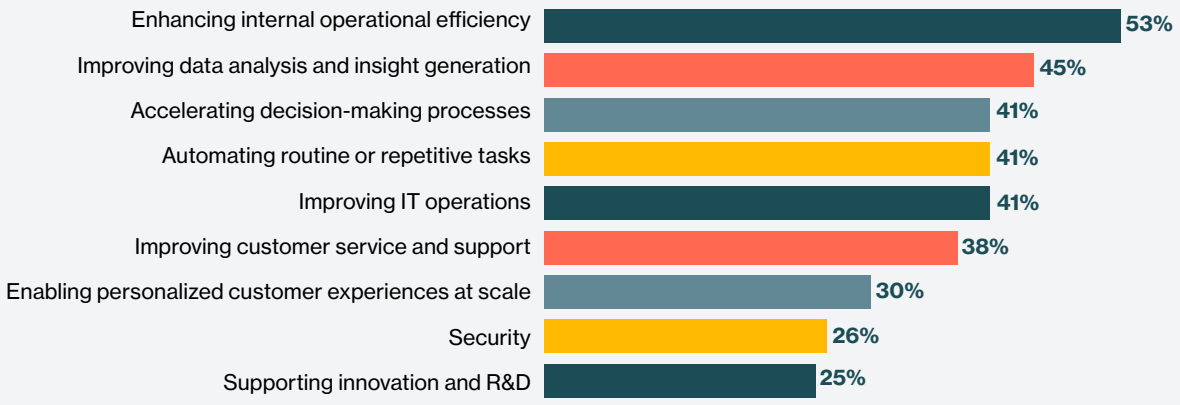

Source: MIT Technology Review Insights survey, 2025

our jobs easier," he says. "It's all internally focused, so nothing outwardly focused where the potential of hallucination could result in giving the wrong decision to a credit union member."

Challenges of agentic AI

According to the experts we interviewed, the chief difficulties associated with implementing agentic AI are not vastly different from those faced with earlier AI generations. The survey respondents agree, pointing in particular to the challenge of integrating it into existing workflows and systems (cited by 47%). Many (38%) also cite the challenge of establishing clear governance and security, and around one-third (34%) point to regulatory and compliance risks (see Figure 11). "Agentic AI demands well-structured, high-context, explainable features — which means that lineage, governance, and semantic clarity become even more critical," says Lee at KT.

For Workday, says Johnson, the main agentic challenges are twofold. "One is authentication, confirming that the agent has the authority to perform an action based on the requester. But the bigger challenge is ensuring that we surface enough data for an agent to understand the actions it should be taking."

Finally, there is the ever-present challenge of improving data quality. "As with AI generally, it's 'garbage in, garbage out,'" says Chaddock. "Again, it all starts with the data quality."

> "Agentic AI demands well-structured, high-context, explainable features — which means that lineage, governance, and semantic clarity become even more critical."
>
> Sejung Lee, Chief Data Officer, KT

Figure 11: Integrating into existing workflows, along with governance and security concerns, are the main difficulties for organizations with agentic AI
(% of respondents)

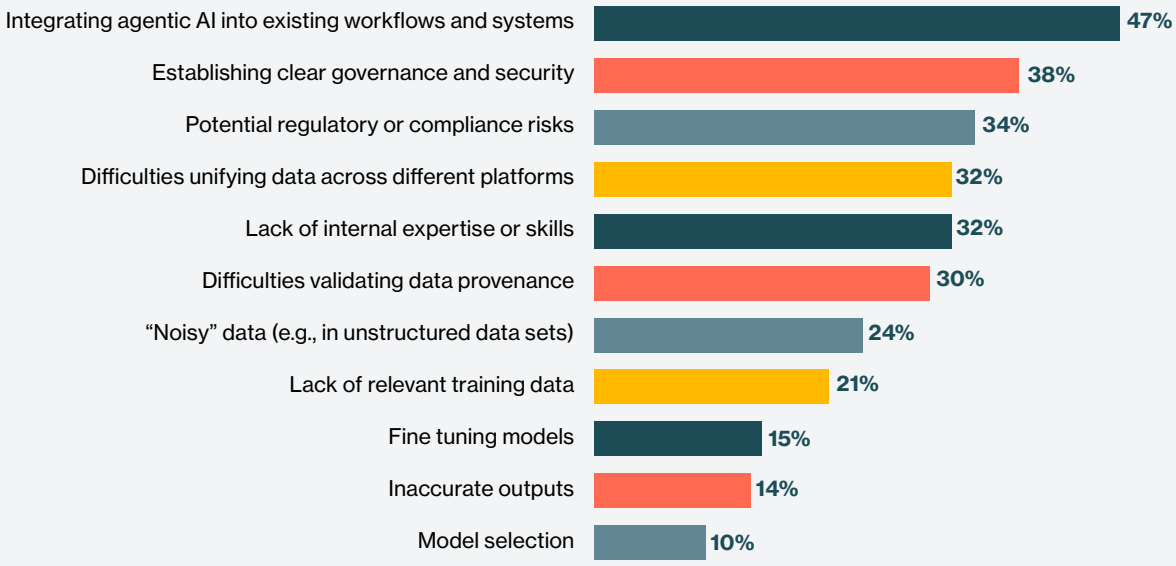

Source: MIT Technology Review Insights survey, 2025

Testing the waters with agentic AI

The organizations highlighted in our study are mainly in exploration mode with agentic AI. There are good reasons for moving slowly, as the survey indicates, not least the challenges associated with integration and governance (see Figure 11). But several businesses are actively testing the waters.

"We're currently prioritizing high-feasibility, low-data, and low-risk use cases," says Rani Johnson of Workday. "For example, if we see that someone's device is not performing well, an agent can be prompted to alert the person [and offer to diagnose and address the reasons]." At the same time, Workday has use cases that Johnson's team needs to understand better. "If someone is taking maternity leave, for instance, they may want to learn details about their benefits. That's deep into data that is very sensitive, and there's going to be a significantly higher threshold to move that use case forward."

The Fonterra Co-operative Group, a dairy business with headquarters in New Zealand, has built a multi-agent solution that is available to over 3,000 of the group's users, says Helius Guimaraes, the company's CDAIO. "It can answer employee queries about the group's leave policy, for instance. It can also tell someone what is impacting the delivery of Fonterra products on time to North America. The agent goes to the database, finds the answer, and returns it. The level of simplification is a big change for us." However, says Guimaraes, his team is not yet able to use these agents to execute tasks. "That's our next step. We're testing it, but those use cases aren't live yet."

As these examples highlight, there are risk equations to be considered when enabling agents to take action. The degree of risk, of course, differs according to the use case, says Irfan Khan, president and chief product officer, SAP Data and Analytics. "For example, there is a high risk association in using agents to help make business decisions that have direct top line or bottom line consequences. You don't want your profit and loss statement to be reflecting incorrect assumptions because your agents didn't quite get the math right. A harmonized data layer is the first critical step towards any successful agentic AI strategy."

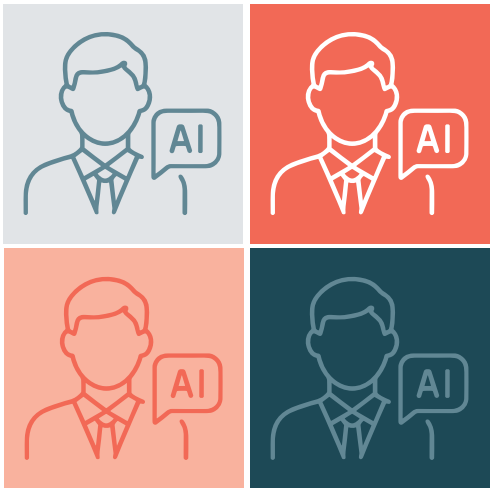

> "There is a high risk association in using agents to help make business decisions that have direct top line or bottom line consequences. You don't want your profit and loss statement to be reflecting incorrect assumptions because your agents didn't quite get the math right."
>
> Irfan Khan, President and Chief Product Officer, SAP Data and Analytics

04 AI-powered data management

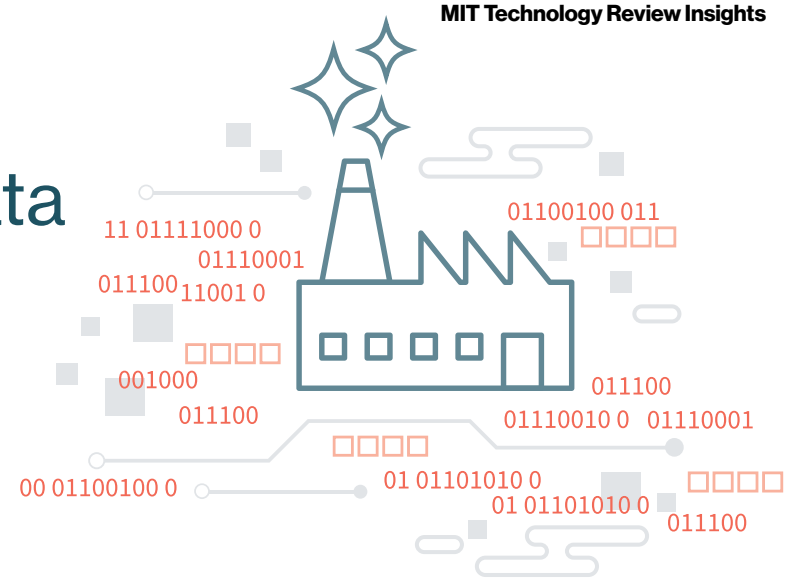

Data teams are increasingly focused on supporting AI models and use cases, especially since the advent of generative AI. Now AI is returning the favor. New agents and tools powered by AI are coming on stream that promise to make engineers, scientists, and architects more productive through intelligent automation. AI code generation tools have come into wide use by engineers as well as application developers. Other types of tools, some supported by agentic AI models, are available to help automate activities such as data cleansing, integration, cataloging, orchestration, and pipeline monitoring.

More than two-thirds of the survey respondents (67%) say their organizations are using AI-powered data management tools today. The rest intend to begin doing so within the next one to two years (see Figure 12). And over half the respondents (54%) say the emergence of such tools has been a reason to update their data strategy.

Figure 12: Current status of data technologies and practices in respondent organizations
(% of respondents)

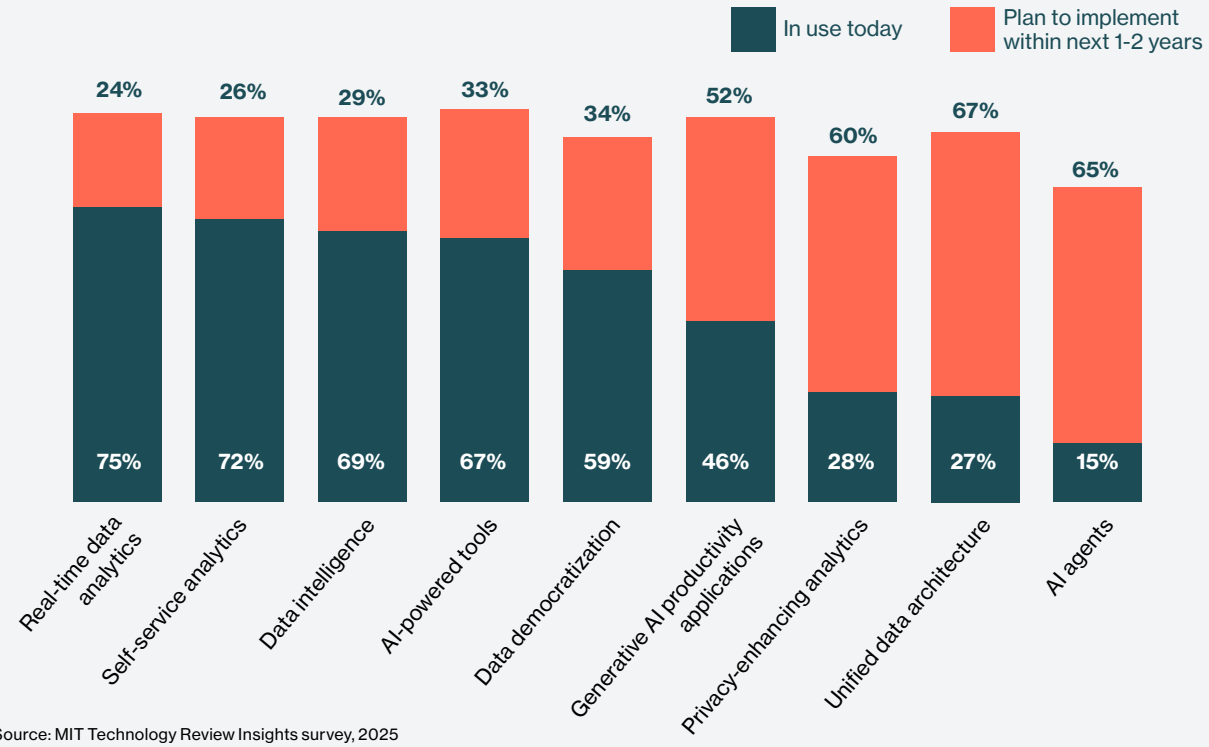

Source: MIT Technology Review Insights survey, 2025

The organizations highlighted in our study are using such tools for a variety of purposes. A key use case at the Navy Federal Credit Union is data cleansing. "AI is helping us to focus on the data we really need to clean," says Chaddock. "In the past, we would go through a long list in a spreadsheet and try to clean everything up. Now we can pre-reduce that list to areas where the models disagree." And Ramachandran says 3M is starting to use machine learning to drive data management. "For example, we're looking at AI tools for data quality management and data observability," he says.

According to Parizot, Reckitt is starting to use AI tools in areas such as data transformation but not yet on a large scale. He expects that eventually to change, however. "The tools are still rudimentary today compared to what I'd think will be possible in the next couple of years," he says.

Data intelligence

Data teams are also using AI to help them gain a clearer picture of their data estates. Nearly seven in ten (69%) of the survey respondents say their organizations use data intelligence today; most of the rest will start using it in the next one to two years (see Figure 12).

Data intelligence is a set of practices involving the analysis of metadata to gain maximum visibility of the organization's data and an understanding of its quality, how it is used, and the impact of its use. "Moving forward, we're focusing more and more on analysis of our metadata," says Expedia Group's Pickrell. "Whether that's our physical data table, metadata of our data lakes, or data structures, it's all vitally important."

With the help of data intelligence, respondents say their organizations mainly aim to improve data governance (cited by 51%). But many also see it helping them to connect data silos, improve how they organize data, and more generally to analyze more of their data (all cited by 40%).

"What makes data more intelligent is adding context to it," says Ramachandran. "That comes through analysis of metadata. A lot of what we are focused on is just descriptions of the data, but also information on what kind of business process the data is involved in and what the value is of decisions made in that process."

"When you're training an AI model, you need data that has business context, and data intelligence brings that," says Khan. "Without the context – knowing where the data is sourced from and its value, standalone and across all your different stakeholders – you won't reap much benefit from it."

"What makes data more intelligent is adding context to it. That comes through analysis of metadata. A lot of what we are focused on is just descriptions of the data, but also information on what kind of business process the data is involved in and what the value is of decisions made in that process."

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

Avoiding agent chaos

Among the many barriers cited by respondents to scaling generative AI, around one fifth (22%) (see Figure 8) – and over one-quarter (27%) of the largest organizations in the survey – cite toolchain fragmentation as a paint point for them.

As AI agents and tools come onto the market from vendors, data and AI teams need to ensure discipline in spending decisions to avoid sprawl and fragmentation. "The third-party landscape is offering more and more native agents," notes Christopher d'Arcy of E.ON. "My nightmare scenario is that we're flooded with uncoordinated agents of all kinds, many of which are possibly redundant."

Murali Vridhachalam of Gilead Sciences harbors the same concern. "We don't want sprawls of agents everywhere," he says. For his organization, using an AI agent marketplace to source such tools makes sense because such platforms typically vet agents for security and performance. "Our team members can discover, re-use, and share agents more reliably using a marketplace than trying to do this on their own." That approach, says Vridhachalam, helps ensure discipline among employees when selecting agents.

For d'Arcy, it is important to create an environment that securely facilitates the interoperability of agents, especially those that may eventually act autonomously. "How is a reasoning layer going to decide which agent to activate in the execution of a task and which ones should be used for which purpose?" he asks. "Failure to orchestrate this is a risk, and we need to get our heads around it early enough in the game."

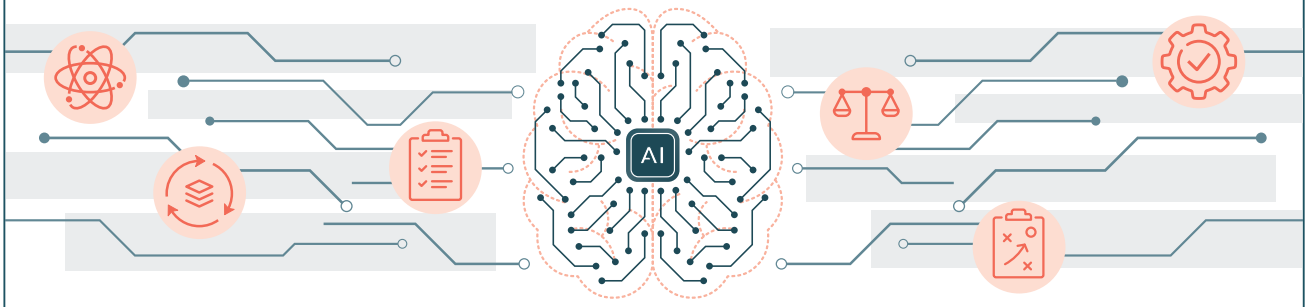

"How is a reasoning layer in a model going to decide which agent to activate in the execution of a task and which ones should be used for which purpose? Failure to orchestrate this is a risk, and we need to get our heads around it early enough in the game."

Christopher d'Arcy, Chief Data & AI Officer and Managing Director, E.ON Digital Technology

05 Freedom to act

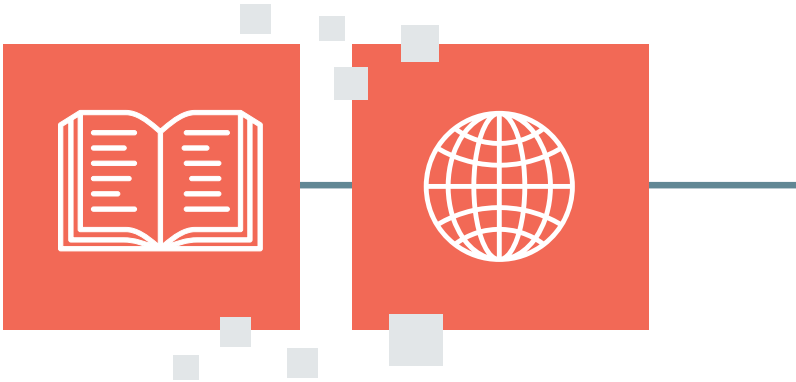

The surge of generative AI in 2022 and 2023 caught many enterprises off guard. Wanting to avoid missing a unique technology opportunity, many committed early to using proprietary generative AI developed by big technology vendors. Such models conferred significant power and performance advantages, but they also limited organizations' ability to experiment and customize.

As AI marches on and new capabilities emerge, the technology leaders we spoke to are determined to keep their model options open. "One of our top priorities has been to ensure platform optionality, not getting locked into one provider," says LeRoy. W.W. Grainger often uses a relatively small, low-cost model for some types of work early in a use case's development and then a more powerful (and expensive) one at a later stage, with another one used when nearing production. "We need the ability to arbitrage across models," says LeRoy. "We don't want to wait until the market matures to pick our partners. We want to be taking advantage right now."

> "One of our top priorities has been to ensure platform optionality, not getting locked into one provider. We need the ability to arbitrage across models."
>
> Jonny LeRoy, Senior Vice President and Chief Technology Officer, W.W. Grainger

For businesses like FOX that generate large volumes of original content, optionality also includes the possibility of building their own generative AI models. "There are a thousand or more models to choose from," says Hildebrandt. "We base our decisions on the opportunities for differentiation. If they exist, it makes sense to put our AI talent on building the best models to leverage our content. It's not something we can get out of the box from a provider."

Optionality for many organizations also includes the possibility of using vendors' proprietary models for specific use cases. "There is a variety of proprietary models we can orchestrate," says Hildebrandt, "and we work with the best of all of them."

Fox Corporation, like most of the other organizations in our study, leans toward the use of open-source models and platforms to the extent possible, says Hildebrandt. "We are model agnostic, but I think open source creates a healthy pressure on the proprietary ecosystem that's to the benefit of everyone."

Cloud flexibility

Optionality extends to organizations' use of cloud infrastructure. One-quarter of the survey respondents say implementing a multi-cloud strategy for data and AI is a

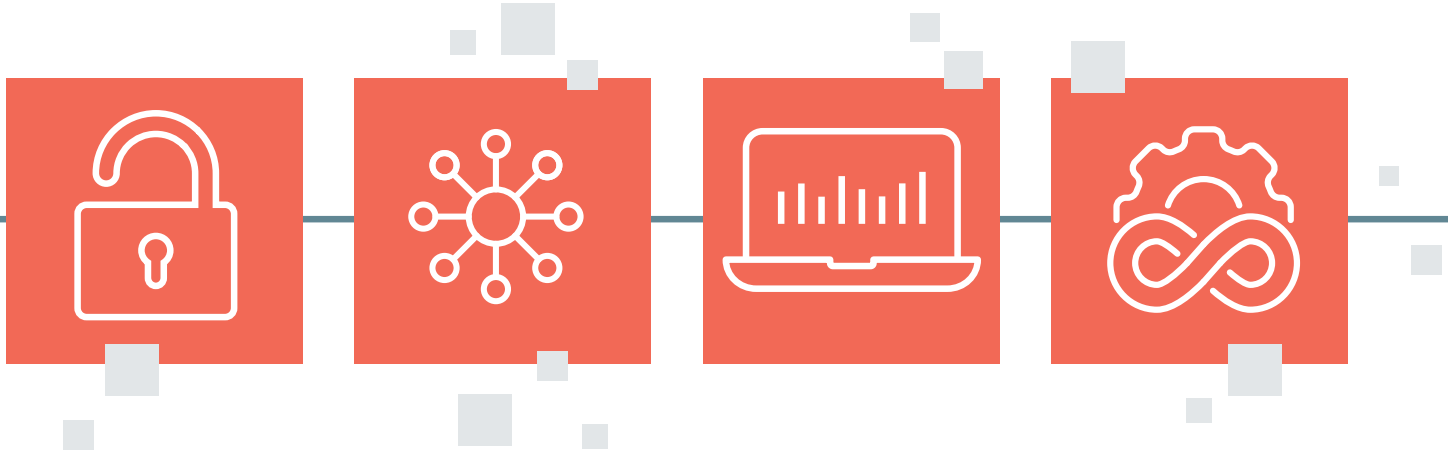

top priority for their organization (see Figure 5). A reason that percentage is not larger may be because many large organizations have been following such an approach for some time.

The Fonterra Co-operative Group has been putting a growing number of data and AI workloads into the cloud in recent years, says Helius Guimaraes. "Doing so means we can worry less about back-up and upgrades and can focus on taking advantage of vendors' new features," he says. "We also bring in data from several different platforms to help us optimize our value chain and run our business more efficiently. We now have unified governance across data and AI within all the different cloud platforms that we use."

Pursuing a multi-cloud strategy is a top priority for 41% of telecom organizations – the highest share of any industry in the survey. And when putting workloads into the cloud, many look for the security advantages that managing them on-premises can provide.

Sejung Lee says KT's approach is cloud-first, "but not cloud at any cost." She explains: "We're building our data and AI infrastructure in secured public cloud environments, with architectural safeguards to ensure that data sovereignty and organizational control are preserved – even in fully managed cloud settings. This foundation, strengthened further by applying responsible AI practices, is what allows us to scale AI with confidence – not just faster, but smarter and safer."

"[Putting data and AI workloads into the cloud] means we can worry less about back-up and upgrades and can focus on taking advantage of vendors' new features. We also bring in data from several different platforms to help us optimize our value chain and run our business more efficiently. We now have unified governance across data and AI within all the different cloud platforms that we use."

Helius Guimaraes, Chief Data & AI Officer, Fonterra Co-operative Group

06 Conclusion

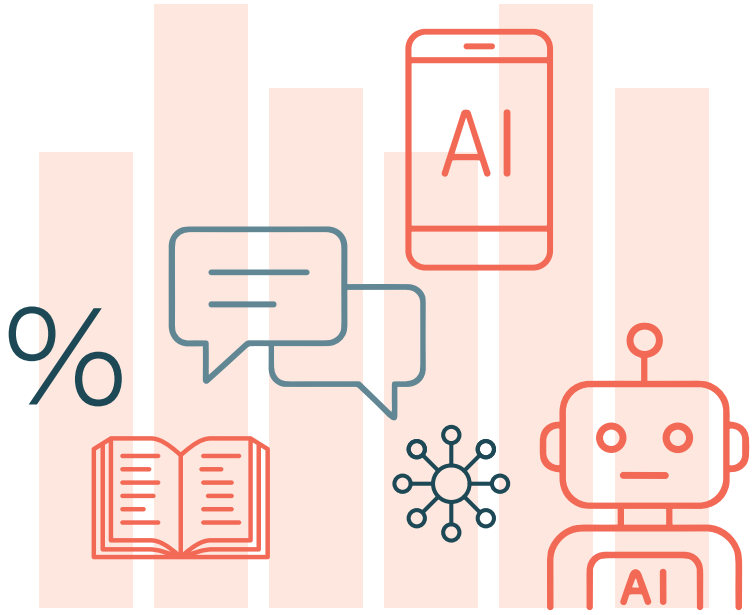

None of the executives we interviewed believe that AI's progression will decelerate any time soon. Just as data and AI teams are becoming familiar with the workings of agentic and multi-modal AI, new advances are on the horizon that will offer more opportunities – and pose new challenges – to teams. Examples include multi-modal AI models 2.0, which can analyze and reason across a wide range of data formats in a unified way; large world models that allow AI to interact with complex physical environments; explainable AI – a set of methods and techniques that make models and their decisions more transparent; and, further into the future, quantum AI, which will greatly increase the computational power of models.

It is difficult as yet to predict the precise nature of the challenges these or other AI advances will pose. But data and AI leaders and their teams can draw on their experiences with generative AI to help devise plans to address them. Four high-level lessons emerge from our research:

Just as data and AI teams become familiar with the workings of agentic and multi-modal AI, new advances are on the horizon that offer more opportunities – and challenges.

Exercise discipline. In the early days of generative AI adoption, many organizations let use case exploration grow freely without tight oversight. The leaders we spoke to for this study are taking a more rigorous approach with agentic AI development, insisting that model inputs and outputs are clearly understood before moving beyond the initial stages.

Keep your options open. AI is moving very fast in terms of models, providers, and products. Not all will stand the test of time. Unlike with the first generative AI models, there are now numerous options, including open-source models, for retaining maximum flexibility and avoiding lock-in.

Avoid needless fragmentation. The world suddenly seems full of agents with AI-based capabilities for automating data management. They offer considerable potential for efficiency gains. But there is also potential for duplication, sprawl, and added complexity. Data teams must exercise discipline in selecting their new tools.

Keep the focus on business outcomes, not technology. Generative AI has shown how senior executives and line employees across the organization can get excited by leaps in AI's capabilities. But too few organizations are generating valuable business outcomes. Just 2% of respondents rate their organization's AI performance highly today in terms of delivering measurable business results from them. Do not do AI for the sake of doing AI.

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
All Illustrations provided by Adobe stock, assembled by Scott Shultz Design.

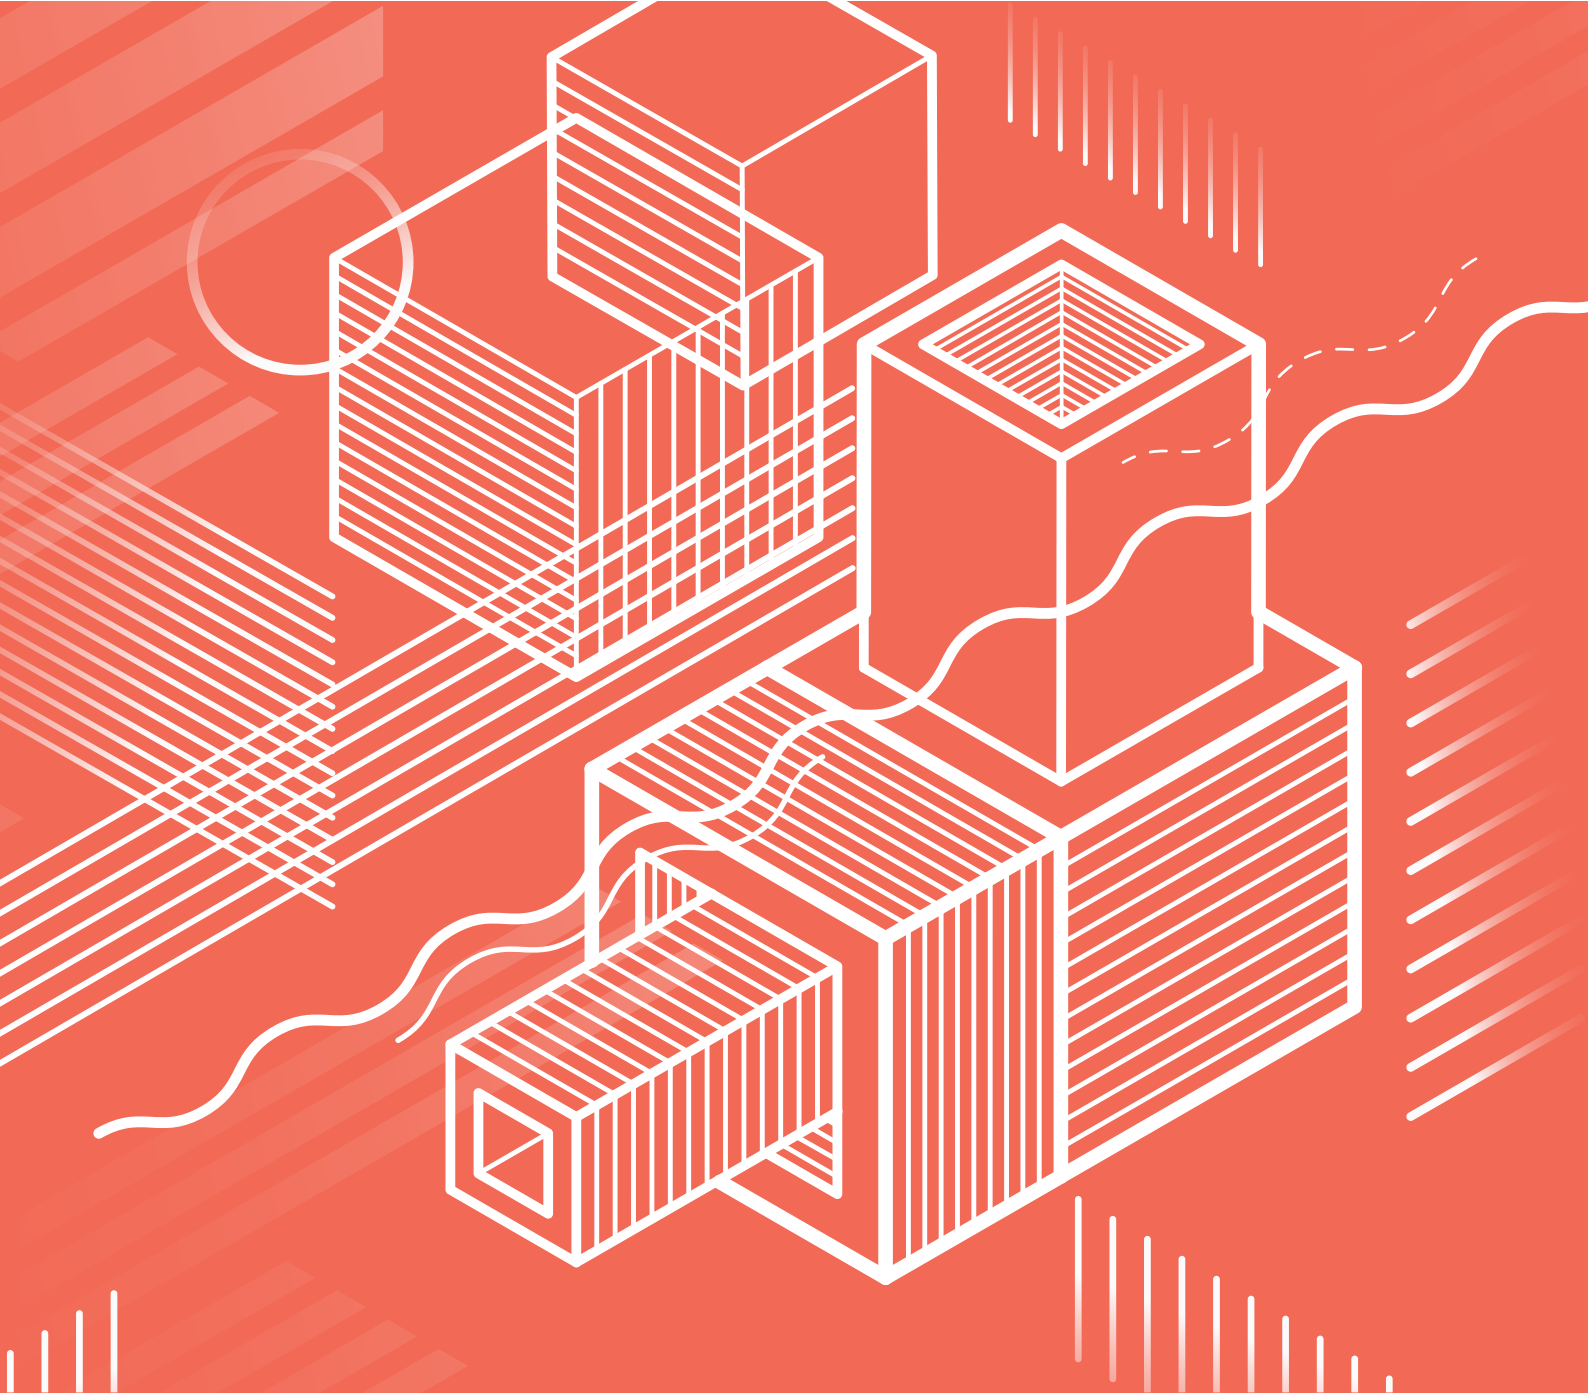

MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com