MIT Technology Review Insights

Produced in partnership with

Microsoft

A ranking of 101 agent tasks reveals where workflows are trending and where connected intelligence is critical.

Agent confidence on the technical frontier

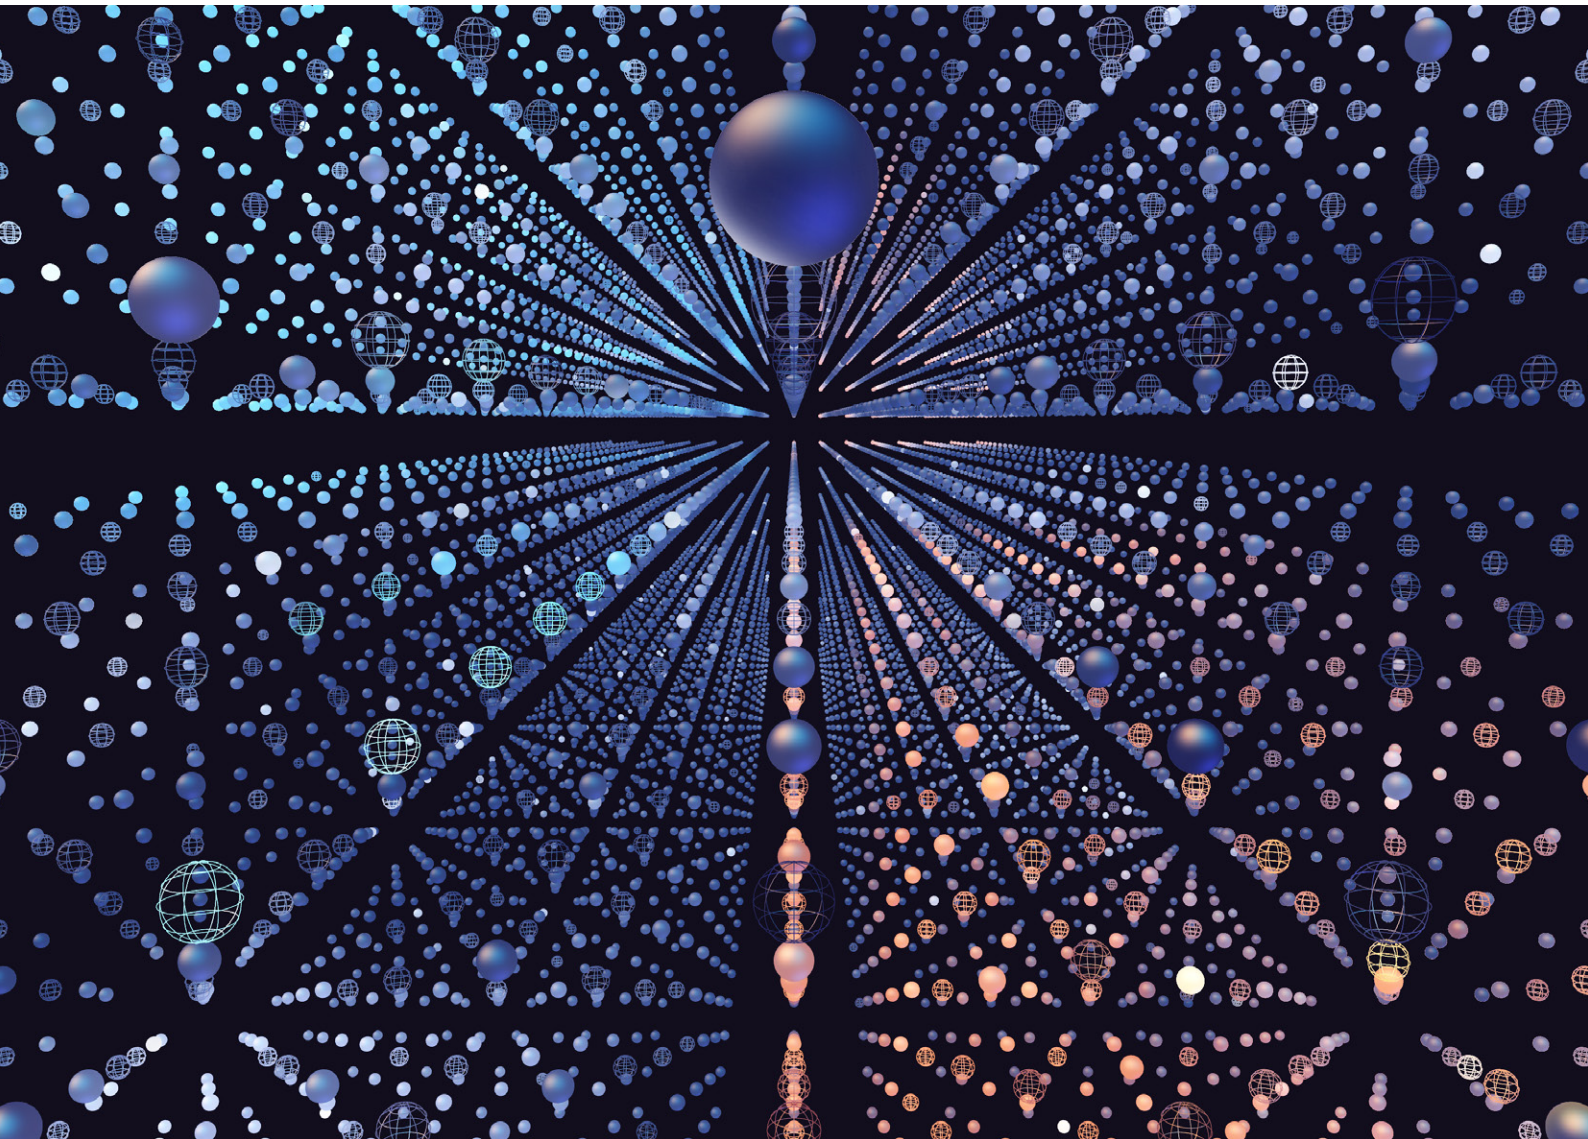

Preface

“Agent confidence on the technical frontier” is an MIT Technology Review Insights report sponsored by Microsoft. Based on survey research and executive interviews, the report maps where technical experts are most confident in agents and how fast transformation is happening.

Our research aims to show how organizations use agents to enable transformation by building agentic workflows for AI, data, and cloud that not only work alongside human teams but also enable tech teams to advance their careers. As pressure increases on executives to deliver immediate, measurable ROI, organizations cannot afford a trial-and-error approach to automation; they need to scale agentic AI safely, empower teams, and provide the context to do so to succeed in an agentic world.

Denis McCauley was the author of the report, Laurel Ruma 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 perspectives:

Amanda Silver, Corporate Vice President, Microsoft 365 Core and Work IQ, Microsoft

Kim Manis, Corporate Vice President of Product, Microsoft Fabric Platform

Jeremy Winter, Corporate Vice President and Chief Product Officer, Microsoft Azure Platform

Contents

01 Executive summary

Enterprise investment in AI is booming. Gartner is calling 2026 an "inflection year" for organizations to align their AI projects with strategic business objectives.[1] As the pressure to prove ROI mounts, executives and technology leaders are looking to agentic AI to drive the measurable financial outcomes their businesses seek.

A prime opportunity for AI agents exists in the tech function, where IT infrastructure costs are projected to grow two to three times by 2030, even as budgets remain unchanged, according to McKinsey.[2] And in the last 18 months, tech teams – the engineers, developers, architects, and other practitioners who are building, deploying, and continually improving their organizations' infrastructure and applications – are clearly putting agents to work.

The ultimate promise of agents is not only to automate tasks but to manage and coordinate entire workflows, pursuing business goals in a way that allows humans and agents to work together. Given the risks involved in automated decision-making, teams cannot delegate the work that agents do without confidence that they are fully capable of performing the task and that it will do so in a safe, reliable, and secure manner.

Among technology experts, our research shows that teams are exceedingly confident about using agentic AI across a significant amount of AI, data, and cloud tasks.

Where agent readiness drops is largely due to a lack of business context being supplied to agentic systems. The more complex the task, the more reasoning capability an agent requires and the greater its need for business context. Such context-generation capabilities for agents are still at an early stage of development, especially in situations where enterprise data is difficult to wrangle and connect into the agent lifecycle at the speed and quality in which developers and executives need it. Human oversight is a key factor of success in deploying agentic AI.

Knowing that tech teams are in a pivotal position to lead this transformation, the experts we interviewed expect agent confidence to accelerate as experience with agents deepens and business environments mature. "As we design agents to operate within the same operational boundaries, identity systems, and governance models that teams already use, they start to behave more like the systems organizations already trust," says Jeremy Winter, corporate vice president and chief product officer at Microsoft Azure Platform.

This report, based on a survey of 300 global technology experts, ranks 101 tasks across AI, data, and cloud workflows based on respondents' confidence in agents acting on their behalf. It also examines how technology teams view the opportunities and challenges related to agentic AI, along with the potential for the technology to enhance their careers.

Key findings from the report include:

- **Confidence in agents is surging for measurable tasks and growing in areas of complex judgment**. Technology experts overwhelmingly believe agents help with everyday work including streamlining processes, improving performance, and reducing repetitive tasks. Confidence is highest for processes

like generating reports and boilerplate code, and there is clear opportunity where tasks involve multi-step workflows and advanced reasoning to make decisions.

- **Data workflows are the breakthrough domain.** Tech teams trust agents most where structure can provide a reliable foundation for decisions. This includes areas such as data quality monitoring, visualization anomaly detection, real-time data stream monitoring, and data profiling. This is where domain experts closest to the point of data generation can provide context to allow agents to act and deliver trusted outcomes.

- **Agents are increasing human capability.** Keeping humans in the loop is critical to 59% of respondents – especially when it comes to the emerging tasks on our list. These complex tasks require supervision, orchestration, and reasoning where human judgement and AI systems work in a continuous loop, each making the other more capable over time. Experts we interviewed believe agents will get to this point faster through human collaboration and the support of increasingly connected business intelligence and well-integrated systems.

- **Systems thinking makes tech experts indispensable.** As with any transformation, the ability to reason across complexity, connect moving parts, and design for outcomes at scale must be developed to empower agents and humans to work together. And since most respondents say the success of agents relies on keeping humans in the loop, it is vital for organizations to invest in upskilling for this level of innovation.

Methodology

In February and March 2026, MIT Technology Review Insights conducted a global survey of 300 technology executives, directors and team leaders, and contributors, to measure the readiness in deploying agents across 101 tasks in AI, data, and cloud workflows. The survey respondents are evenly distributed among those with domain expertise in AI and machine learning, data management, and cloud and IT management:

- **AI workflows:** AI development, machine learning, coding or similar
- **Data workflows:** Data development, analytics or similar
- **Cloud workflows:** Cloud, IT management, operations, or similar

The respondents work in organizations across 12 industries, including:

- Consumer goods and retail
- Financial services
- Hospitality, travel, and leisure
- Electricity and energy production
- IT, telco, and technology
- Manufacturing
- Media and marketing
- Pharmaceuticals
- Health care and life sciences
- Professional services
- Government and public services
- Education

The respondents work in organizations ranging from startups to companies earning upward of $10 billion in revenue annually.

For the ranking of 101 AI agent tasks, the scores are derived from the confidence ratings provided by the 300 survey respondents on a zero to 100 scale. The higher the score, the greater the survey respondents' confidence in AI agents to perform the task. To ensure the relevance of their responses, practitioners with specific domain expertise rated only tasks within their domain.

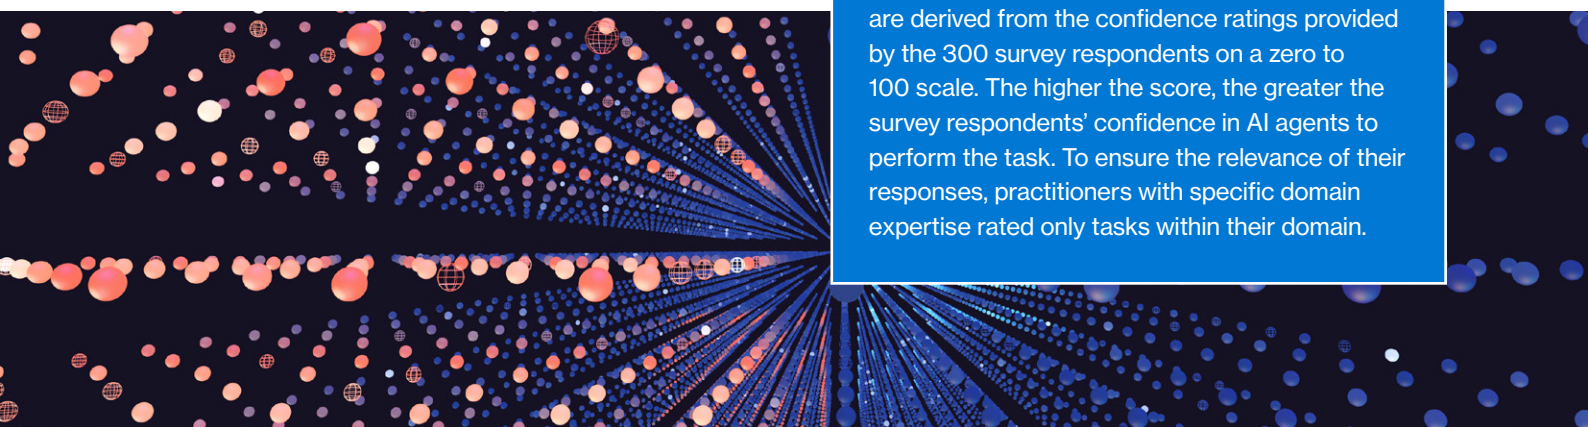

02

Measuring confidence in agentic AI use cases

While most tech teams understand that agentic AI will eventually transform how they work, finding the best way to use agents today can be murky. "Teams need to understand what agents can actually do for them now," says Amanda Silver, corporate vice president, Microsoft 365 Core and Work IQ at Microsoft. "I think there are many use cases where coding agents can actually help teams with common problems like technical debt, for example—it's an opportunity cost for every team."

"Generally you want to apply agents when you have decisions that are reversible, when you have scoped, understandable work. That's exactly the kind of use case that agents are really great for."

To provide technology teams with guidance in this area, we scored 101 specific AI, data, and cloud tasks based on their current readiness for agent use. The scores are derived from the confidence ratings provided by the 300 survey respondents on a zero to 100 scale. The higher the score, the greater the survey respondents' confidence in AI agents to perform the task.

Trending agent use cases

The trending tasks in our agent confidence ranking are straightforward ones that developers may find tedious and are often happy to hand over to agents, largely involving content generation and coding. "automated generation of business reports and their distribution to stakeholders" has the highest confidence score at 83.5, while "boilerplate code generation for new software features" follows closely behind at 82.5 (see Figure 1).

Confidence is everything when it comes to agents, says Winter. "The biggest blocker to agent adoption isn't its capability. It's tech teams' confidence in the agent and its ability to execute accurately and responsibly inside the operational workflows."

Generally speaking, the more straightforward and low-risk the task is, the more confidence practitioners have in agents to execute it reliably. "Agents are great at executing very scoped, understandable work," says Silver. "Tasks that involve more ambiguity and high-stakes decisions, or where convincing data is lacking, are those we should still direct to humans."

The 2026 Agent Confidence Index with the entirety of the 101 tasks ranked can be found in the last chapter of this report.

"What we thought agents were capable of just a few months ago has already changed profoundly."

Amanda Silver, Vice President, Microsoft 365 Core and Work IQ, Microsoft

Figure 1: AI agents show immediate value with straightforward and measurable tasks
(all respondents)

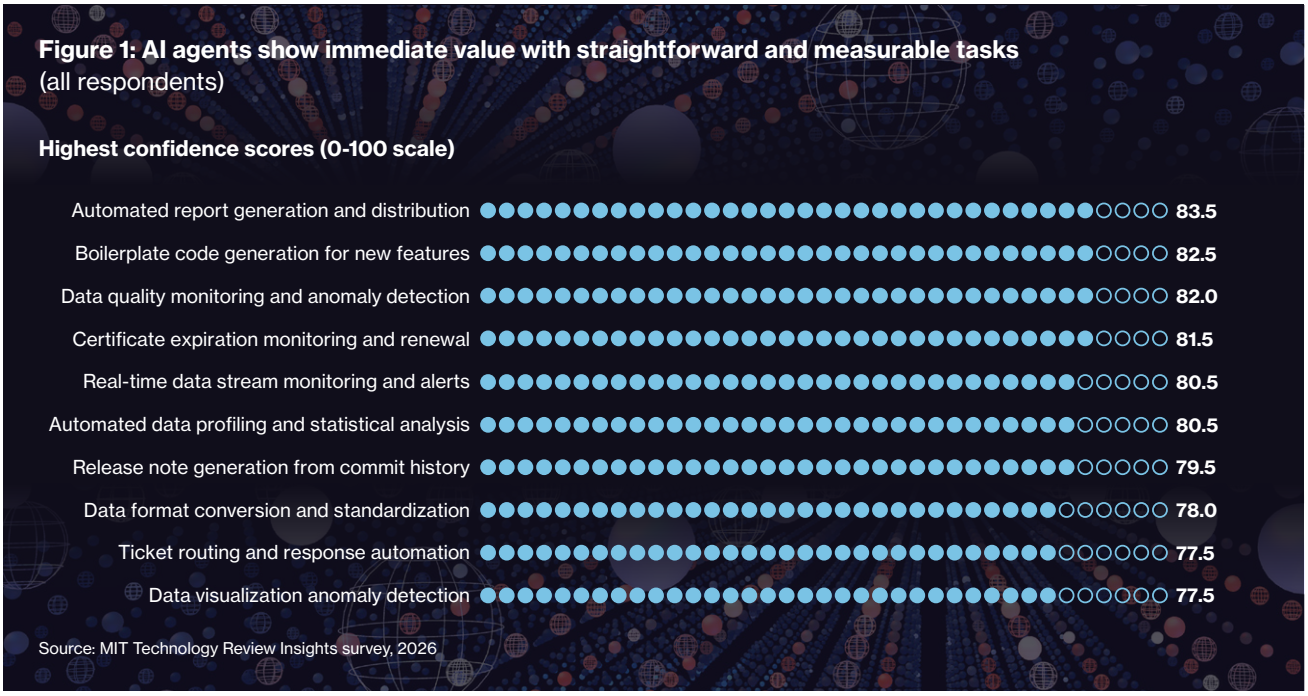

"These are some of the first killer use cases for agents," adds Silver. "One reason trust is high in agent use for boilerplate code is because we can reliably measure its efficacy. For example, we can look at merge rates to the main code base to see if the generated code is high enough quality to be merged in."

Winter says he is seeing organizations build confidence over time: "The teams seeing the most value from agents are starting to narrow—clearly scoped work, clear checkpoints, clear ownership. When teams see agents handling these tasks cleanly over time, confidence builds and allows them to expand into more complex, multi-step workflows across systems." The goal is to transform the business in what Bain Consulting calls a "structural shift in enterprise technology."[3]

Data-specific tasks also dominate what's trending, including data quality monitoring, real-time data stream monitoring, and data profiling. "AI has removed so much toil that used to be involved in working with data," explains Kim Manis, corporate vice president of product at Microsoft Fabric Platform. "That's why there's so much positivity toward using data and AI together. You see the gains right away. The key is the sooner you start, the better the gains will be."

"AI has removed so much toil that used to be involved in working with data. That's why there's so much positivity toward using data and AI together. You see the gains right away."

Kim Manis, Corporate Vice President of Product, Microsoft Fabric Platform

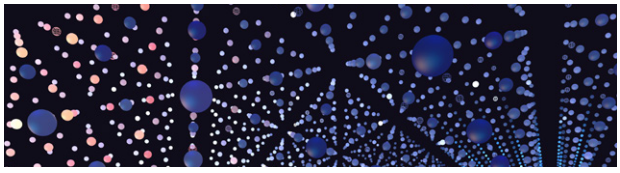

Figure 2: Confidence is building for AI agents handling complex tasks (all respondents)

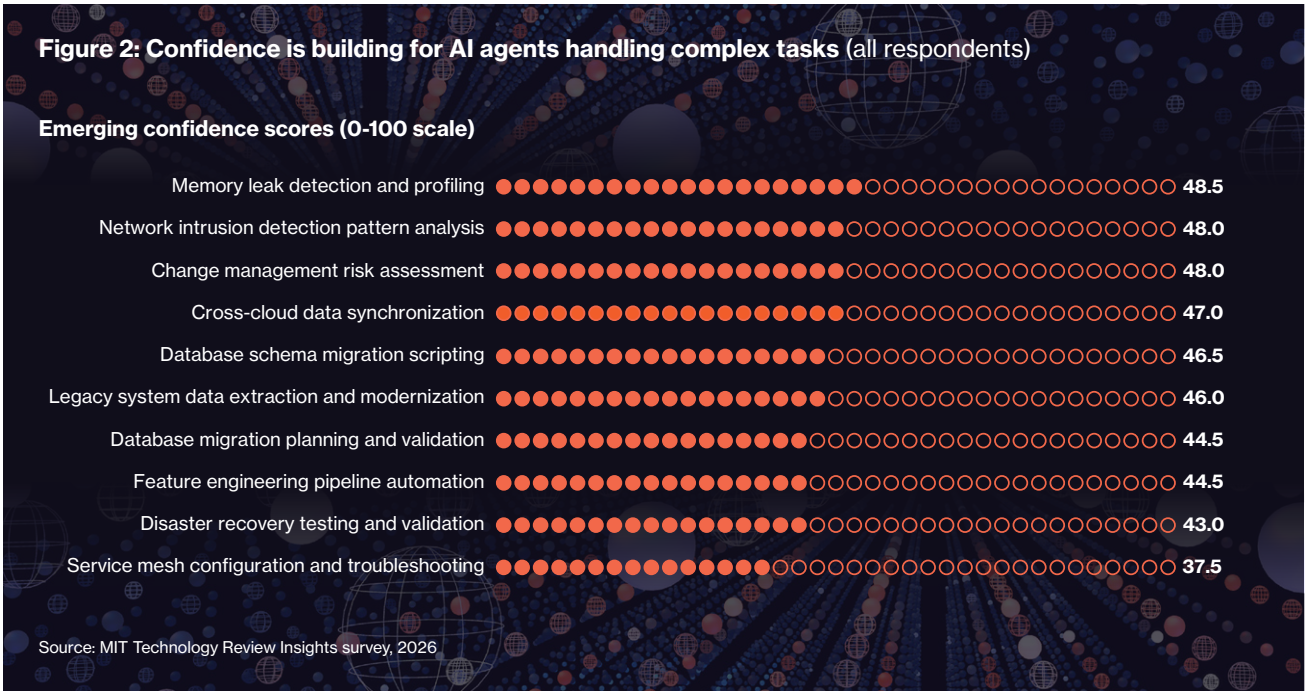

Source: MIT Technology Review Insights survey, 2026

Context is king

These agent confidence scores show there is an opportunity to address a critical business need – explaining business context to AI. Generally speaking, the more complex the task, the more an agent needs to have the right business context – to be connected to the necessary and correct data and workflows, and have an intelligence layer that can tie it all together. Take the examples of disaster recovery testing and database migration planning. These are not only multi-step operations, they also require considerable business information to make accurate decisions. This could be information about past experiences, existing relationships, interactions between different business units, or about broader organizational objectives.

As agents become more autonomous, the quality of their decisions depends increasingly on the quality of the context they are given. Which is why it is critical to give agents access to the right business data, workflows, and organizational knowledge to confidently execute and avoid costly rework.

Not many technology teams are currently able to ensure agents have such context, says Kim Manis, corporate vice president of product at Microsoft Fabric Platform. “We see so many scenarios where customers say an agent is getting the decisions wrong, but it turns out that the agent just doesn’t have the context.”

And context is needed even for some seemingly straightforward tasks, says Manis. “If you hire a new data analyst and ask them to list the top 10 customers by revenue for the last year, they will need to know which customer column to use, how the company calculates revenue, or which calendar to use, fiscal year or calendar year. An agent needs the same context.”

Figure 3: Trending and emerging agent tasks for data workflows (all respondents)

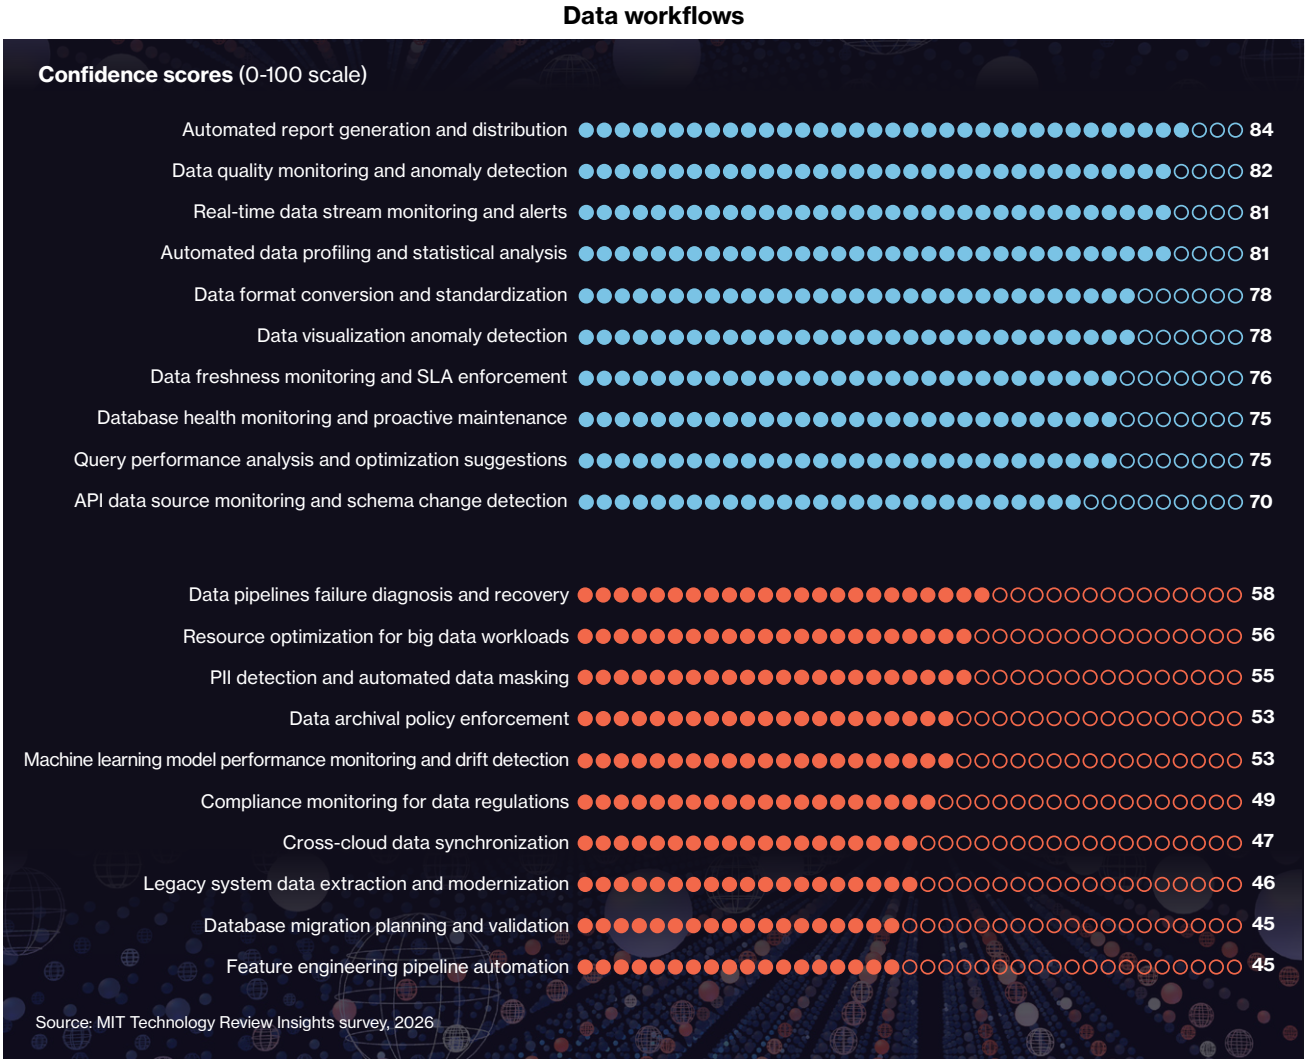

Emerging agent use cases

At the other end of the spectrum are the tasks respondents deem less ready for agents today (see Figure 2). These emerging use cases consist of some of the most complex technology tasks. They may involve the coordination of long, multi-step workflows as with database migration and planning, for example. "These are complicated tasks even for humans," says Manis. "But we can put in the right scaffolding and infrastructure for teams to make those trade-offs."

The level of task complexity also depends on how interconnected it is with processes elsewhere. Disaster recovery is a good example, says Winter. "Agents just can't operate in isolation; they're touching real infrastructure – policies and systems that other things depend on. Actions taken in areas like disaster recovery can affect availability, resiliency, and application behavior across the environment."

Emerging use cases also involve the orchestration of multiple agents. And several require advanced reasoning capabilities in order to make decisions and take actions. This requires memory and understanding of business context that AI models are only just beginning to develop.

Figure 4: Trending and emerging agent tasks for AI workflows (all respondents)

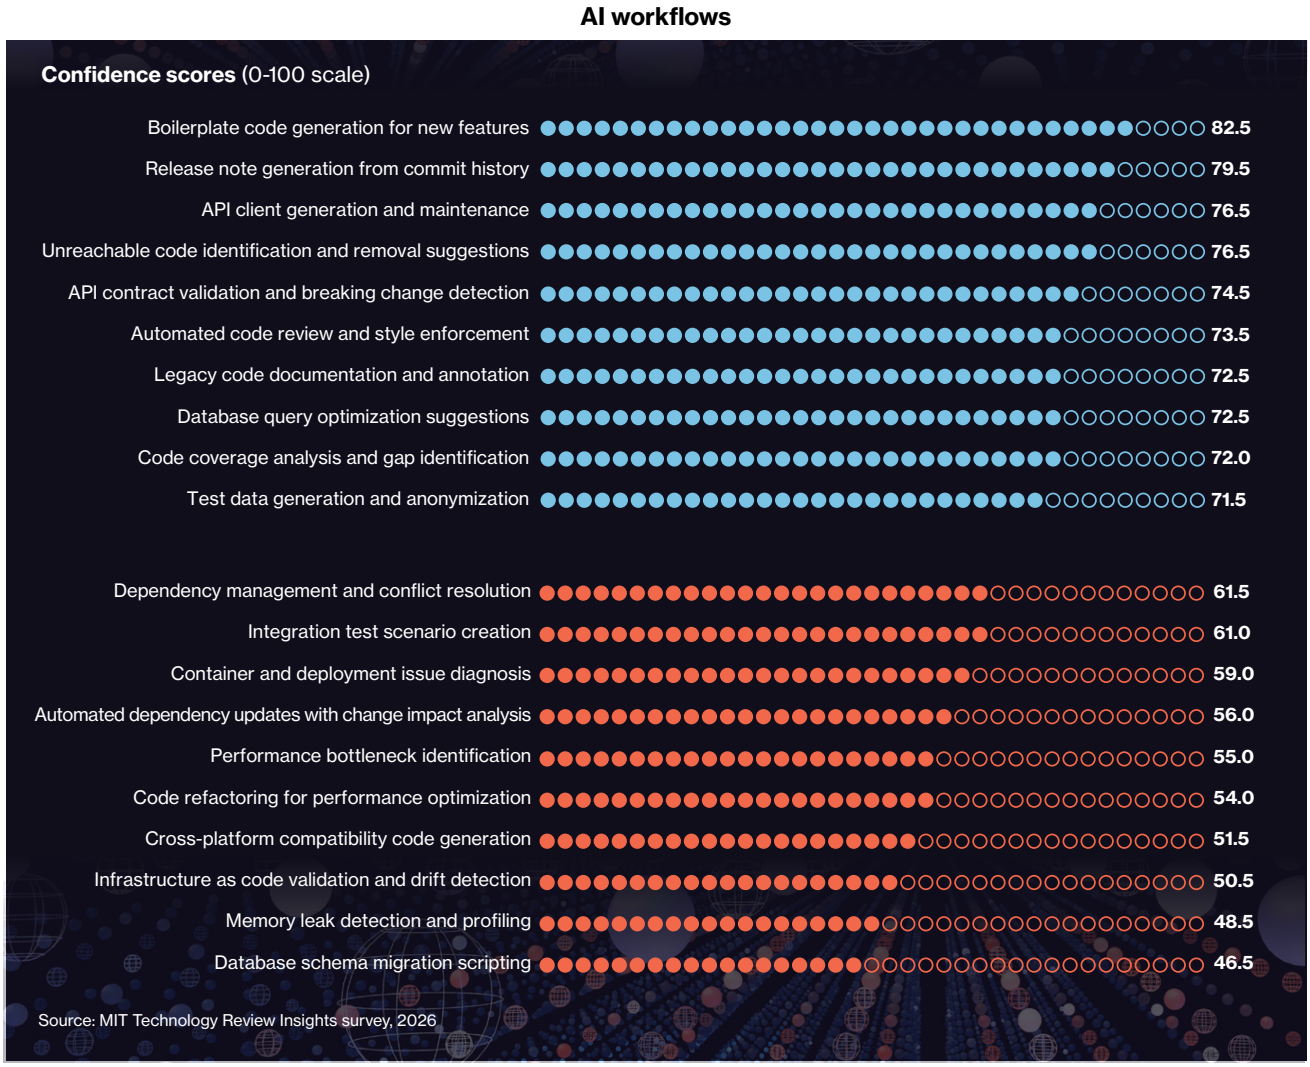

Human engineers and architects are likely to remain heavily involved in such workflows for at least the next couple of years, even while agents execute individual tasks within them. But the experts we interviewed foresee agents becoming increasingly capable of managing entire workflows. “The technology behind agents designed for database schema migration and memory leak detection is advancing quickly on both fronts,” says Silver. “What we thought agents were capable of just a few months ago has already changed profoundly.”

Likewise, says Winter, the confidence gap in areas such as disaster recovery is going to close over time. “We’re already seeing people become more confident with agent use as they iterate,” he says. “As the confidence grows, agent autonomy can progressively expand.”

Figure 5: Trending and emerging agent tasks for cloud workflows (all respondents)

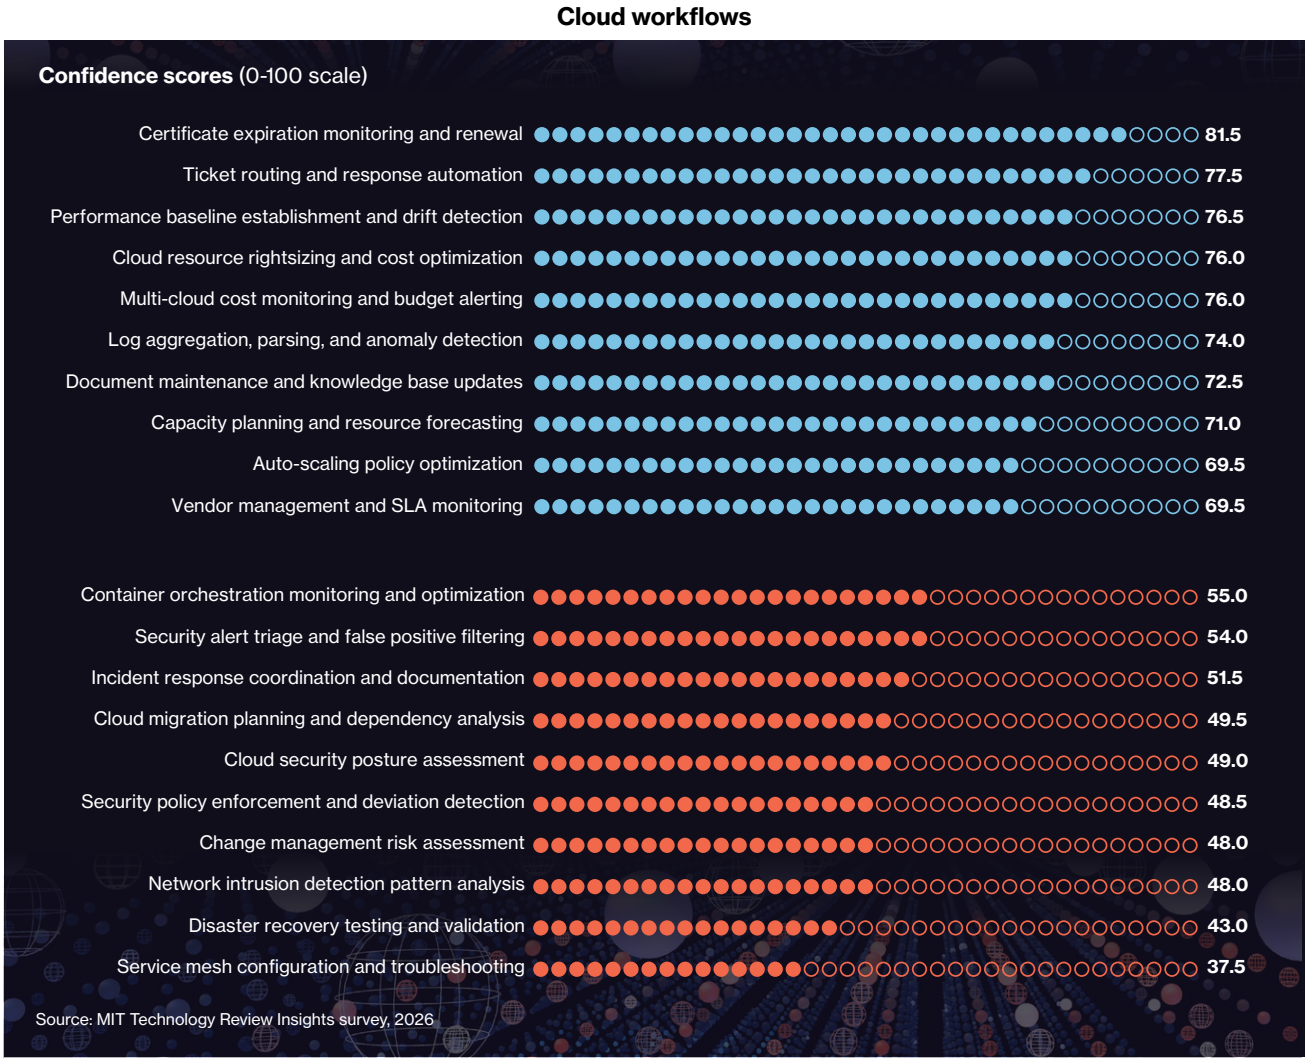

"Agents just can't operate in isolation; they're touching real infrastructure – policies and systems that other things depend on."

Jeremy Winter, Corporate Vice President and Chief Product Officer, Microsoft Azure Platform

03 Transforming workflows with agents

The trending tasks for agents are relatively straightforward, and low risk, but they also align with the areas of impact that technology experts are most likely to seek. When asked about the opportunities that agents offer for improving everyday work, the survey respondents point first and foremost to improved efficiency, with half (51%) citing streamlining processes (see Figure 6).

"The real opportunity with agents isn't automating simple tasks. It's simplifying long, complex workflows that require heavy coordination. Cloud migration is a good example here. Teams spend a significant amount of time stitching together context across systems, approvals, and dependencies. Agents can really change the equation here," explains Winter.

The importance attached to some of these goals differs considerably depending on the respondents' respective roles. Executives, for example, are more likely than team leads or contributors to focus on how agents can enable scale. Business growth, after all, is an overarching goal for executives, and AI-powered automation can provide the technology foundations for that. Team leads are focused mainly on streamlining their teams' processes with the help of agents. Contributors, meanwhile, put almost equal emphasis on agent-driven improvements in process efficiency, performance, and reduced repetition.

Figure 6: Efficiency leads for agent use (top responses, overall and by role)

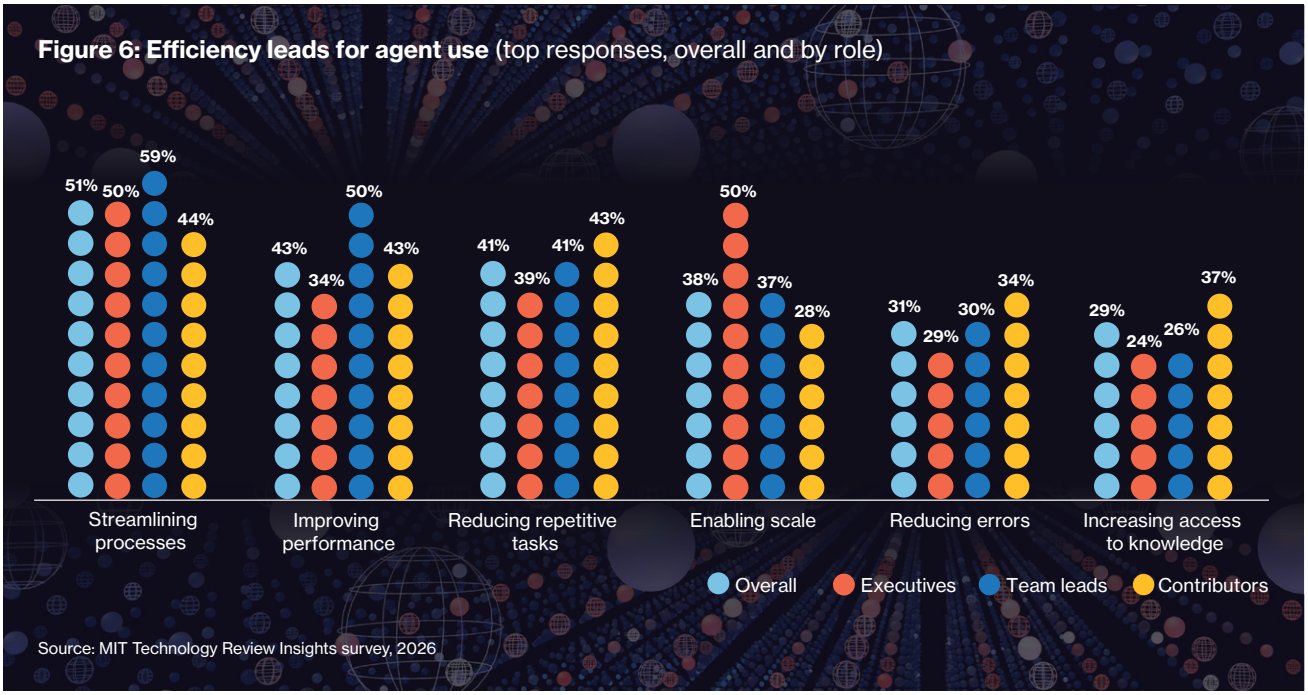

Source: MIT Technology Review Insights survey, 2026

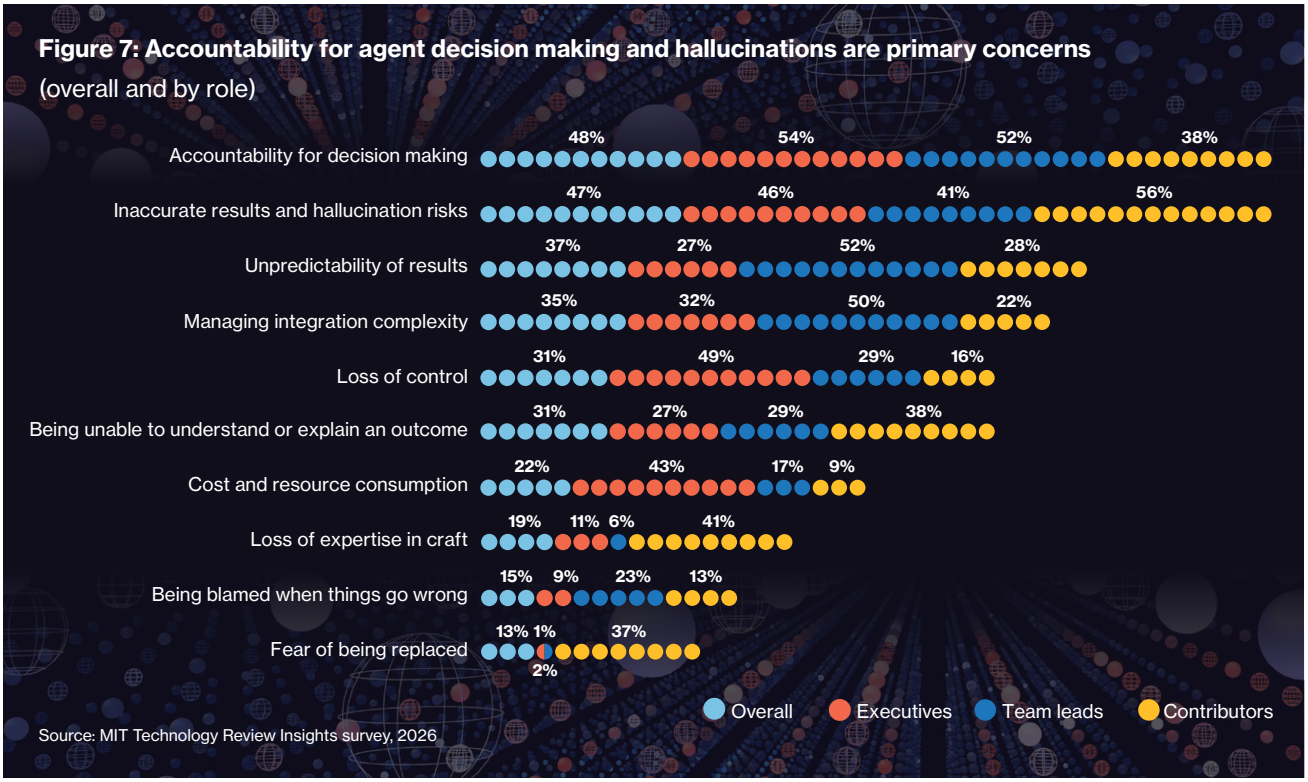

Where agents need guidance

Two main concerns surfaced for tech teams. First, accountability for decisions made by agents (cited by 48% overall) and, second, the potential for inaccuracy and hallucination in their outputs (47%) (see Figure 7).

Here, too, the respondents' role in their organization influences their responses. Focused on practical utility, individual contributors are primarily concerned with hallucinations (56%), as well as the potential loss of expertise in their craft (41%), where as executives are focused on accountability concerns (54%). After all, executives and team leads are more likely than contributors to have to answer if agent decisions are proven wrong. Team leads are equally concerned about the unpredictability of results. This may relate to the adaptive and probabilistic decision-making of the more advanced agents coming into use.

To address such concerns, most respondents say they plan to keep humans in the loop (cited by 59% overall). Over half (53%) are also monitoring agent activity closely and tracing their decision inputs as a means of increasing observability (see Figure 8). "With scenarios that are high stakes or irreversible, it is humans that need to make the key decisions," says Silver. "Agents can recommend, but ultimately humans must have the final say."

"With scenarios that are high stakes or irreversible, it is humans that need to make the key decisions. Agents can recommend, but ultimately humans must have the final say."

Amanda Silver, Corporate Vice President, Microsoft 365 Core and Work IQ, Microsoft

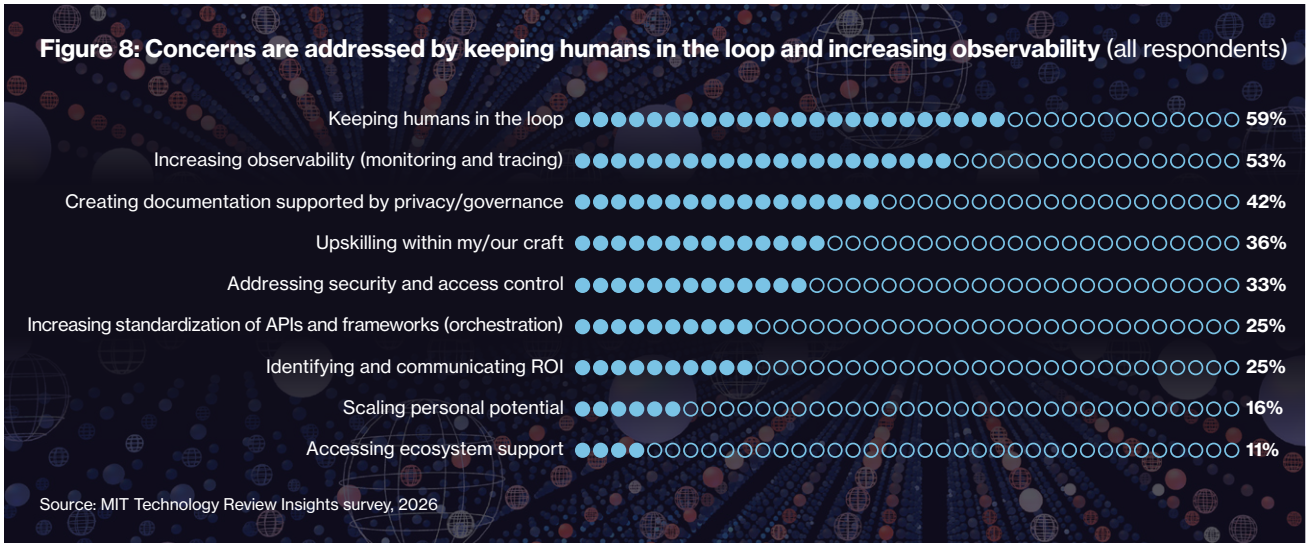

The expansion of agentic AI poses another risk—to the very profession of software development and engineering. Many contributors (41%) worry about losing some of the expertise they have in their craft as agents take over more tasks. Some (37%) also worry they will be replaced by agents.

In a recent publication of Communications of the ACM, Mark Russinovich, CTO, deputy CISO, and technical fellow of Microsoft Azure, and Scott Hanselman, vice president and member of technical staff at Microsoft, wrote that such expertise loss could harm the prospects not only of junior team members but also of tech teams and their businesses more widely.[4] "Organizations will need skilled developers and engineers to get the most out of agentic AI," say Russinovich and Hanselman. "But the talent pipeline will dwindle if teams stop hiring junior staff on the back of automation. There will then be no path in which young practitioners can develop the ability to steer, verify, and integrate AI output."

Organizations will need to help junior and more experienced tech members alike become accustomed to working with agents.

Looking through the industry lens

In our survey results we found some interesting industry data bubble up for retail, manufacturing, financial services, and health care.

The manufacturing industry is the most concerned (55%) with inaccurate results and hallucination risks, while the other three industries are focused on accountability. Each of the industries are tackling concerns in different ways. Financial services respondents are focused on creating documentation supported by privacy and governance (70%), while health care, the other regulated industry, is focused on increasing observability (76%). Manufacturing most highly ranks the importance of keeping humans in the loop (76%), followed by retail (68%).

As for the trending tasks, manufacturing respondents rank "automated report generation and distribution" as the task they were most confident in, followed by "database health monitoring and proactive maintenance." Retail respondents rank "automated data profiling and statistical analysis" first, followed by "certificate expiration monitoring and renewal." Financial services respondents rank "certificate expiration" first, and then "ticket routing" and "response automation" were tied for second place. Lastly, health care respondents rank both "boilerplate code generation for new features" and "release note generation from commit history" as trending tasks.

04 Systems thinking for an agentic world

A vast majority of technology experts believe agents will help their careers more than harm them. For cloud workflows, 96% of respondents are either “confident” or “very confident” that using agents for system reliability and site operations will help their career prospects (see Figure 9). And 92% of AI workflow respondents feel the same about using agents for evaluation and quality assurance. In fact, in all the areas we explored, more than 75% of respondents in each of their respective workflows say they are either “confident” or “very confident” that agentic AI use will help their careers.

Figure 9: Widespread confidence that AI agents will enhance careers (all respondents)

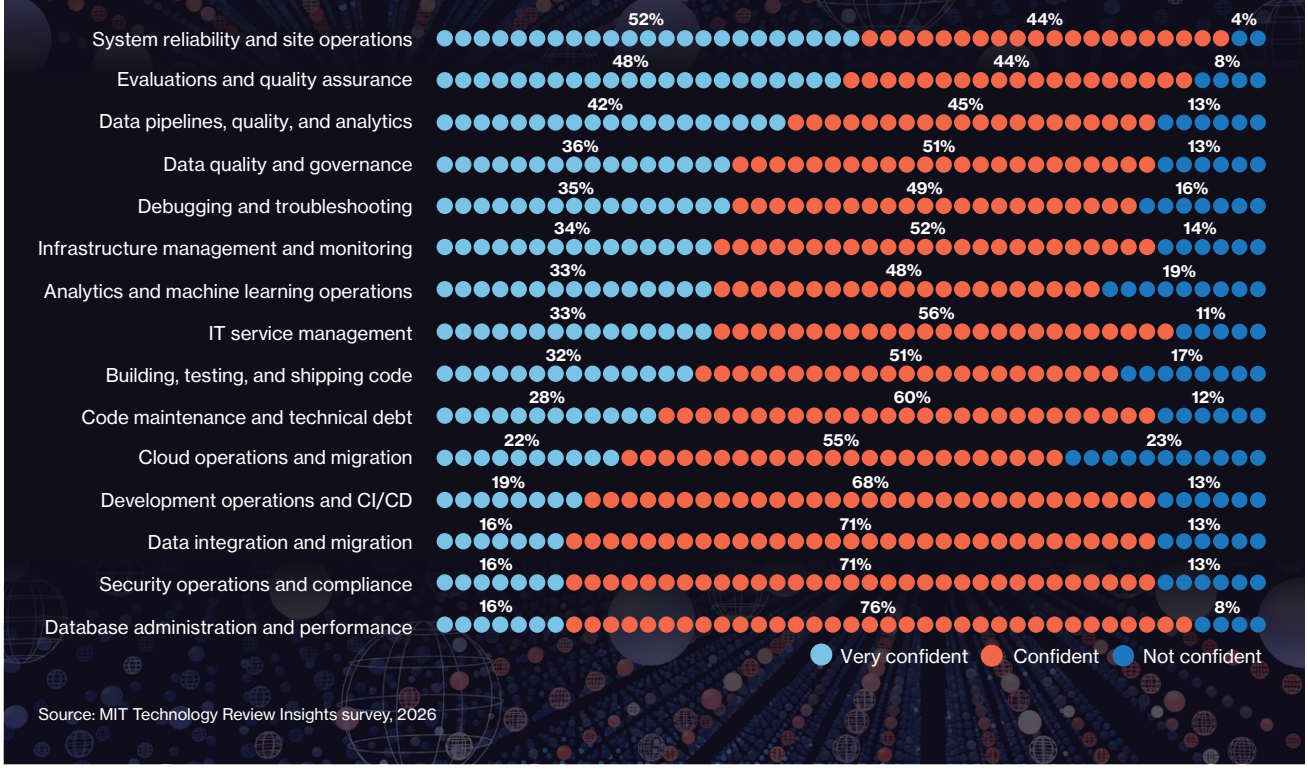

Source: MIT Technology Review Insights survey, 2026

"As agents take on more of these repetitive tasks, engineers can spend more time understanding the systems, diagnosing problems, reviewing decisions, evaluating trade-offs, and defining the intent of where the product or business should go."

Jeremy Winter, Corporate Vice President and Chief Product Officer, Microsoft Azure Platform

Developers and engineers also understand that agent-driven automation can give them time and space to hone their skills in other areas, says Winter. "As agents take on more of these repetitive tasks, engineers can spend more time understanding the systems, diagnosing problems, reviewing decisions, evaluating trade-offs, and defining the intent of where the product or business should go," he says.

"AI can reduce the toilsome processes for team members so that they can get to the interesting work, the fun work, the work that requires their brain power," adds Manis.

Learning from agents and mentors

What types of skills, then, will team members need to make themselves indispensable in an agentic environment? "Systems thinking is crucial in this environment," says Silver. "Building software with multiple constituent parts or coordinating multiple agents that interact together is complex. To manage that and reason about it requires exceptional concentration and skill."

Russinovich and Hanselman agree: "Even the most reliable systems cannot fully replace the judgment, creativity, and adaptability required to handle uncertainty, make complex decisions, and maintain security. Human oversight, critical thinking, and domain knowledge are indispensable for ensuring the reliability and effectiveness of agentic systems."

Working with agents gives junior developers an opportunity to learn these and other skills faster than their predecessors could, believes Manis. "In the past, a junior developer would apprentice a more senior developer, sitting next to them observing the way the latter got their job done. Junior developers now have this opportunity to learn from agents first. We need to shift our expectations of what a newer member on the team can do."

And Winter believes junior staff are coming in well-equipped to adapt quickly to the agentic AI environment. "Early-in-career engineers coming out of university today understand AI and agentic as part of the core operation model that they'll be operating in," he says.

At the same time, learning higher-level systems skills from senior peers will remain critical within teams. "Organizations cannot just assume that early-in-career staff will benefit from the same productivity gains as seniors," say Russinovich and Hanselman. "Organizations must invest in them. That means giving junior members direct exposure to debugging, making design trade-offs, and building systems – the fundamentals needed to critically evaluate AI output."

Russinovich and Hanselman advocate for creating a "preceptor program" within tech teams. "Experienced mentors, or preceptors, can guide junior team members, teaching them how to direct agentic AI tools, develop critical judgment, and learn the production function of senior engineers."

"The future of software engineering will be defined not by the volume of code AI can generate," add Russinovich and Hanselman, "but by how effectively humans learn, reason, and mature alongside these systems." Companies that help early-in-career practitioners develop these capabilities will serve them well for the agentic future.

05 Conclusion

Agentic AI has the potential to transform how organizations of all types manage their technology operations. That ability to manage entire workflows, to learn and adapt, and to take actions semi-independently provides for an entirely different level of automation than any previous enterprise technology. AI can't be "bolted on" to existing environments, it requires new ways of thinking and deployment throughout systems, processes, and workflows, as well as to be integrated into the way humans work on a daily basis.

Enabled tech teams can move away from manual churn. And by embracing agentic AI, these teams can capitalize on parts of their everyday work where they are already making decisions, being creative, and building best-in-class systems, processes, and product for their customers.

As this research makes clear, however, agents will not remove humans from technology. Rather, agentic AI poses distinct new challenges, specifically: adapting to change, evolving processes, and expanding mindsets at operational and management levels.

The first step to meeting those challenges is to view agent use as an opportunity rather than a threat. The practitioners we surveyed appear to have adopted that mindset. Their widespread confidence in AI agents to perform many of our 101 tasks signals that tech teams strongly value relief from repetitive burden and mistakes, among other benefits.

It is that combination of forces – where agents are performing the most core tasks and being overseen by confident technology experts – that will deliver on agentic AI's potential for true business transformation.

"AI can reduce the toilsome processes for team members so that they can get to the interesting work, the fun work, the work that requires their brain power."

Kim Manis, Corporate Vice President of Product, Microsoft Fabric Platform

06

The 2026 Agent Confidence Index

The 2026 Agent Confidence Index is a ranking of 101 agent tasks. The scores are derived from the confidence ratings provided by 300 survey respondents on a zero to 100 scale. The higher the score, the greater the survey respondents' confidence in AI agents to perform the task. To ensure the relevance of their responses, practitioners with specific domain expertise rated only tasks within their domain.

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Data	Analytics and machine learning operations	Automated report generation and distribution	83.5	84.8	80.5	86.5
AI	Building, testing, and shipping code	Boilerplate code generation for new features	82.5	78.3	85.9	82.9
Data	Data quality and governance	Data quality monitoring and anomaly detection	82.0	80.3	82.9	82.7
Cloud	Infrastructure management and monitoring	Certificate expiration monitoring and renewal	81.5	84.7	81.6	76.9
Data	Pipelines, quality, and analytics	Real-time data stream monitoring and alerts	80.5	86.4	76.8	78.8
Data	Data quality and governance	Automated data profiling and statistical analysis	80.5	72.7	81.7	88.5
AI	Development operations and CI/CD	Release note generation from commit history	79.5	83.3	73.4	81.6

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Data	Data integration and migration	Data format conversion and standardization	78.0	66.7	82.9	84.6
Cloud	IT service management	Ticket routing and response automation	77.5	83.3	76.3	71.2
Data	Analytics and machine learning operations	Data visualization anomaly detection	77.5	78.8	78.0	75.0
Cloud	System reliability and site operations	Performance baseline establishment and drift detection	76.5	76.4	75.0	78.8
AI	Building, testing, and shipping code	API client generation and maintenance	76.5	68.3	82.8	77.6
AI	Code maintenance and optimization	Unreachable code identification and removal suggestions	76.5	80.0	73.4	76.3
Cloud	Infrastructure management and monitoring	Cloud resource rightsizing and cost optimization	76.0	73.6	75.0	80.8
Cloud	Cloud operations and migration	Multi-cloud cost monitoring and budget alerting	76.0	76.4	71.1	82.7
Data	Data quality and governance	Data freshness monitoring and SLA enforcement	75.5	81.8	70.7	75.0
Data	Database administration and performance	Database health monitoring and proactive maintenance	75.0	74.2	76.8	73.1
Data	Database administration and performance	Query performance analysis and optimization suggestions	75.0	77.3	69.5	80.8
AI	Debugging and troubleshooting	API contract validation and breaking change detection	74.5	76.7	78.1	69.7

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Cloud	System reliability and site operations	Log aggregation, parsing, and anomaly detection	74.0	72.2	78.9	69.2
AI	Development operations and CI/CD	Automated code review and style enforcement	73.5	78.3	70.3	72.4
AI	Building, testing, and shipping code	Legacy code documentation and annotation	72.5	73.3	71.9	72.4
Cloud	IT service management	Document maintenance and knowledge base updates	72.5	72.2	76.3	67.3
AI	Debugging and troubleshooting	Database query optimization suggestions	72.5	75.0	71.9	71.1
AI	Evaluations and quality assurance	Code coverage analysis and gap identification	72.0	71.7	75.0	69.7
AI	Evaluations and quality assurance	Test data generation and anonymization	71.5	68.3	67.2	77.6
AI	Evaluations and quality assurance	Regression testing automation after code changes	71.5	65.0	79.7	69.7
Cloud	System reliability and site operations	Capacity planning and resource forecasting	71.0	72.2	69.7	71.2
AI	Development operations and CI/CD	Environment consistency validation	70.0	61.7	76.6	71.1
AI	Evaluations and quality assurance	Automated security vulnerability scanning in dependencies	70.0	68.3	70.3	71.1

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Data	Data integration and migration	API data source monitoring and schema change detection	69.5	60.6	74.4	73.1
AI	Code maintenance and optimization	License compliance checking across dependencies	69.5	61.7	71.9	73.7
Cloud	Cloud operations and migration	Auto-scaling policy optimization	69.5	72.2	68.4	67.3
Cloud	IT service management	Vendor management and SLA monitoring	69.5	66.7	69.7	73.1
AI	Debugging and troubleshooting	Cross-browser compatibility and breaking change detection	68.5	63.3	70.3	71.1
Cloud	Security operations and compliance	Compliance reporting and audit preparation	68.0	70.8	67.1	65.4
Data	Pipelines, quality, and analytics	ETL performance optimization and bottleneck identification	68.0	63.6	67.1	75.0
Data	Pipelines, quality, and analytics	Automated data pipeline testing and validating	68.0	69.7	67.1	67.3
Cloud	System reliability and site operations	Late-night incident triage and initial diagnosis	68.0	72.2	67.1	63.5
AI	Evaluations and quality assurance	Unit test generation and edge case coverage	68.0	65.0	68.8	69.7
Data	Database administration and performance	Backup validating and recovery testing	67.5	59.1	73.2	69.2
Cloud	Security operations and compliance	User access review and privilege cleanup	67.5	69.4	64.5	69.2

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Cloud	System reliability and site operations	Health check automation and dependency mapping	67.5	66.7	69.7	65.4
Cloud	Infrastructure management and monitoring	Database maintenance task scheduling and monitoring	67.5	66.7	67.1	69.2
AI	Debugging and troubleshooting	Stack trace analysis and error categorization	67.0	65.0	71.9	64.5
Cloud	Security operations and compliance	Vulnerability assessment coordination and patch prioritization	67.0	68.1	67.1	65.4
Data	Pipelines, quality, and analytics	Data lineage tracking and impact analysis	67.0	66.7	68.3	65.4
Cloud	Cloud operations and migration	Cloud backup validation and recovery testing	66.5	62.5	68.4	69.2
Data	Database administration and performance	Index optimization and unused index cleanup	66.5	60.6	73.2	63.5
Cloud	IT service management	Asset inventory maintenance and lifecycle tracking	66.5	69.4	61.8	69.2
AI	Development operations and CI/CD	Build failure analysis and suggested fixes	66.0	63.3	65.6	68.4
AI	Code maintenance and optimization	Code smell detection and refactoring	66.0	70.0	67.2	61.8
Data	Analytics and machine learning operations	Data warehouse optimization and query acceleration	65.5	65.2	64.6	67.3

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Data	Database administration and performance	Database capacity planning and growth forecasting	65.0	63.6	64.6	67.3
Cloud	Infrastructure management and monitoring	Network change detection and validation	65.0	61.1	69.7	63.5
Cloud	Infrastructure management and monitoring	Multi-environment configuration synchronization	65.0	66.7	64.5	63.5
AI	Development operations and CI/CD	Deployment pipeline optimization	65.0	63.3	65.6	65.8
AI	Evaluations and quality assurance	Bug reproduction and minimal test case creation	64.5	63.3	65.6	64.5
Data	Data quality and governance	Data catalog maintenance and metadata management	64.5	62.1	70.7	57.7
AI	Code maintenance and optimization	Technical debt measurement and prioritization	64.5	56.7	68.8	67.1
AI	Building, testing, and shipping code	Configuration file generation and environment setup	64.0	60.0	75.0	57.9
Cloud	System reliability and site operations	Runbook execution and validation	62.5	59.7	53.9	78.8
Data	Data quality and governance	Cross-system data consistency validation	62.5	66.7	59.8	61.5
AI	Evaluations and quality assurance	Performance test script creation	62.0	60.0	62.5	63.2
Cloud	Infrastructure management and monitoring	Server patching coordination and rollback planning	61.5	61.1	65.8	55.8

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
AI	Development operations and CI/CD	Dependency management and conflict resolution	61.5	65.0	60.9	59.2
AI	Evaluations and quality assurance	Integration test scenario creation	61.0	60.0	62.5	60.5
Data	Analytics and machine learning operations	A/B test statistical analysis and result interpretation	60.5	56.1	64.6	59.6
Cloud	System reliability and site operations	Automated root cause analysis for system outages	60.0	63.9	57.9	57.7
AI	Debugging and troubleshooting	Container and deployment issue diagnosis	59.0	56.7	59.4	60.5
Data	Data integration and migration	Change data capture setup and monitoring	59.0	63.6	56.1	57.7
Data	Pipelines, quality, and analytics	Scheme evolution detection and pipeline adaptation	58.5	54.5	59.8	61.5
Data	Data integration and migration	Data migration validation and reconciliation	57.5	56.1	54.9	63.5
Data	Analytics and machine learning operations	Machine learning experiment tracking and model versioning	57.5	53.0	58.5	61.5
Data	Pipelines, quality, and analytics	Data pipelines failure diagnosis and recovery	57.5	62.1	56.1	53.8
AI	Code maintenance and optimization	Automated dependency updates with change impact analysis	56.0	51.7	53.1	61.8
Data	Pipelines, quality, and analytics	Resource optimization for big data workloads	55.5	47.0	61.0	57.7

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Cloud	Cloud operations and migration	Container orchestration monitoring and optimization	55.0	61.1	55.3	46.2
AI	Debugging and troubleshooting	Performance bottleneck identification	55.0	56.7	59.4	50.0
Data	Data quality and governance	PII detection and automated data masking	54.5	50.0	57.3	55.8
Cloud	Security operations and compliance	Security alert triage and false positive filtering	54.0	59.7	48.7	53.8
AI	Building, testing, and shipping code	Code refactoring for performance optimization	54.0	63.3	53.1	47.4
Data	Database administration and performance	Data archival policy enforcement	53.0	50.0	51.2	59.6
Data	Analytics and machine learning operations	Machine learning model performance monitoring and drift detection	52.5	51.5	57.3	46.2
Cloud	Security operations and compliance	Incident response coordination and documentation	51.5	52.8	46.1	57.7
AI	Building, testing, and shipping code	Cross-platform compatibility code generation	51.5	50.0	53.1	51.3
AI	Development operations and CI/CD	Infrastructure as code validation and drift detection	50.5	45.0	56.3	50.0
Cloud	Cloud operations and migration	Cloud migration planning and dependency analysis	49.5	55.6	43.4	50.0

Domain	Type of task	Task	Score	Contributor/tech function score	Director/team lead score	Executive score
Data	Data quality and governance	Compliance monitoring for data regulations	49.0	47.0	46.3	55.8
Cloud	Cloud operations and migration	Cloud security posture assessment	49.0	45.8	44.7	59.6
Cloud	Security operations and compliance	Security policy enforcement and deviation detection	48.5	45.8	51.3	48.1
AI	Debugging and troubleshooting	Memory leak detection and profiling	48.5	41.7	60.9	43.4
Cloud	Security operations and compliance	Network intrusion detection pattern analysis	48.0	44.4	46.1	55.8
Cloud	IT service management	Change management risk assessment	48.0	44.4	47.4	53.8
Data	Data integration and migration	Cross-cloud data synchronization	47.0	57.6	40.2	44.2
AI	Building, testing, and shipping code	Database schema migration scripting	46.5	46.7	51.6	42.1
Data	Data integration and migration	Legacy system data extraction and modernization	46.0	40.9	46.3	51.9
Data	Database administration and performance	Database migration planning and validation	44.5	39.4	41.5	55.8
Data	Analytics and machine learning operations	Feature engineering pipeline automation	44.5	45.5	42.7	46.2
Cloud	Infrastructure management and monitoring	Disaster recovery testing and validation	43.0	44.4	38.2	48.1
Cloud	Cloud operations and migration	Service mesh configuration and troubleshooting	37.5	31.9	34.2	50.0

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 entirely 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 Microsoft

Microsoft's mission is to empower every person and every organization on the planet to achieve more.

Endnotes:

1. "Gartner Forecasts Worldwide AI Spending to Grow 47% in 2026," Gartner, May 19, 2026, https://www.gartner.com/en/newsroom/press-releases/2026-05-19-gartner-forecasts-worldwideai-spending-to-grow-47-percent-in-2026.
2. "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity," METR, July 10, 2025, https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/.
3. Edd Gent, "AI coding is now everywhere. But not everyone is convinced," MIT Technology Review, December 15, 2025, https://www.technologyreview.com/2025/12/15/1128352/rise-of-ai-coding-developers-2026/.
4. Mark Russinovich and Scott Hanselman, "Redefining the Software Engineering Profession for AI," Communications of the ACM, February 17, 2026, https://dl.acm.org/doi/10.1145/3779312.

Cover image by Green Room Inc. Layout and design by Tim Huxford.

To cite this report, please use: "Agent confidence on the technical frontier," MIT Technology Review Insights and Microsoft, June 2026.

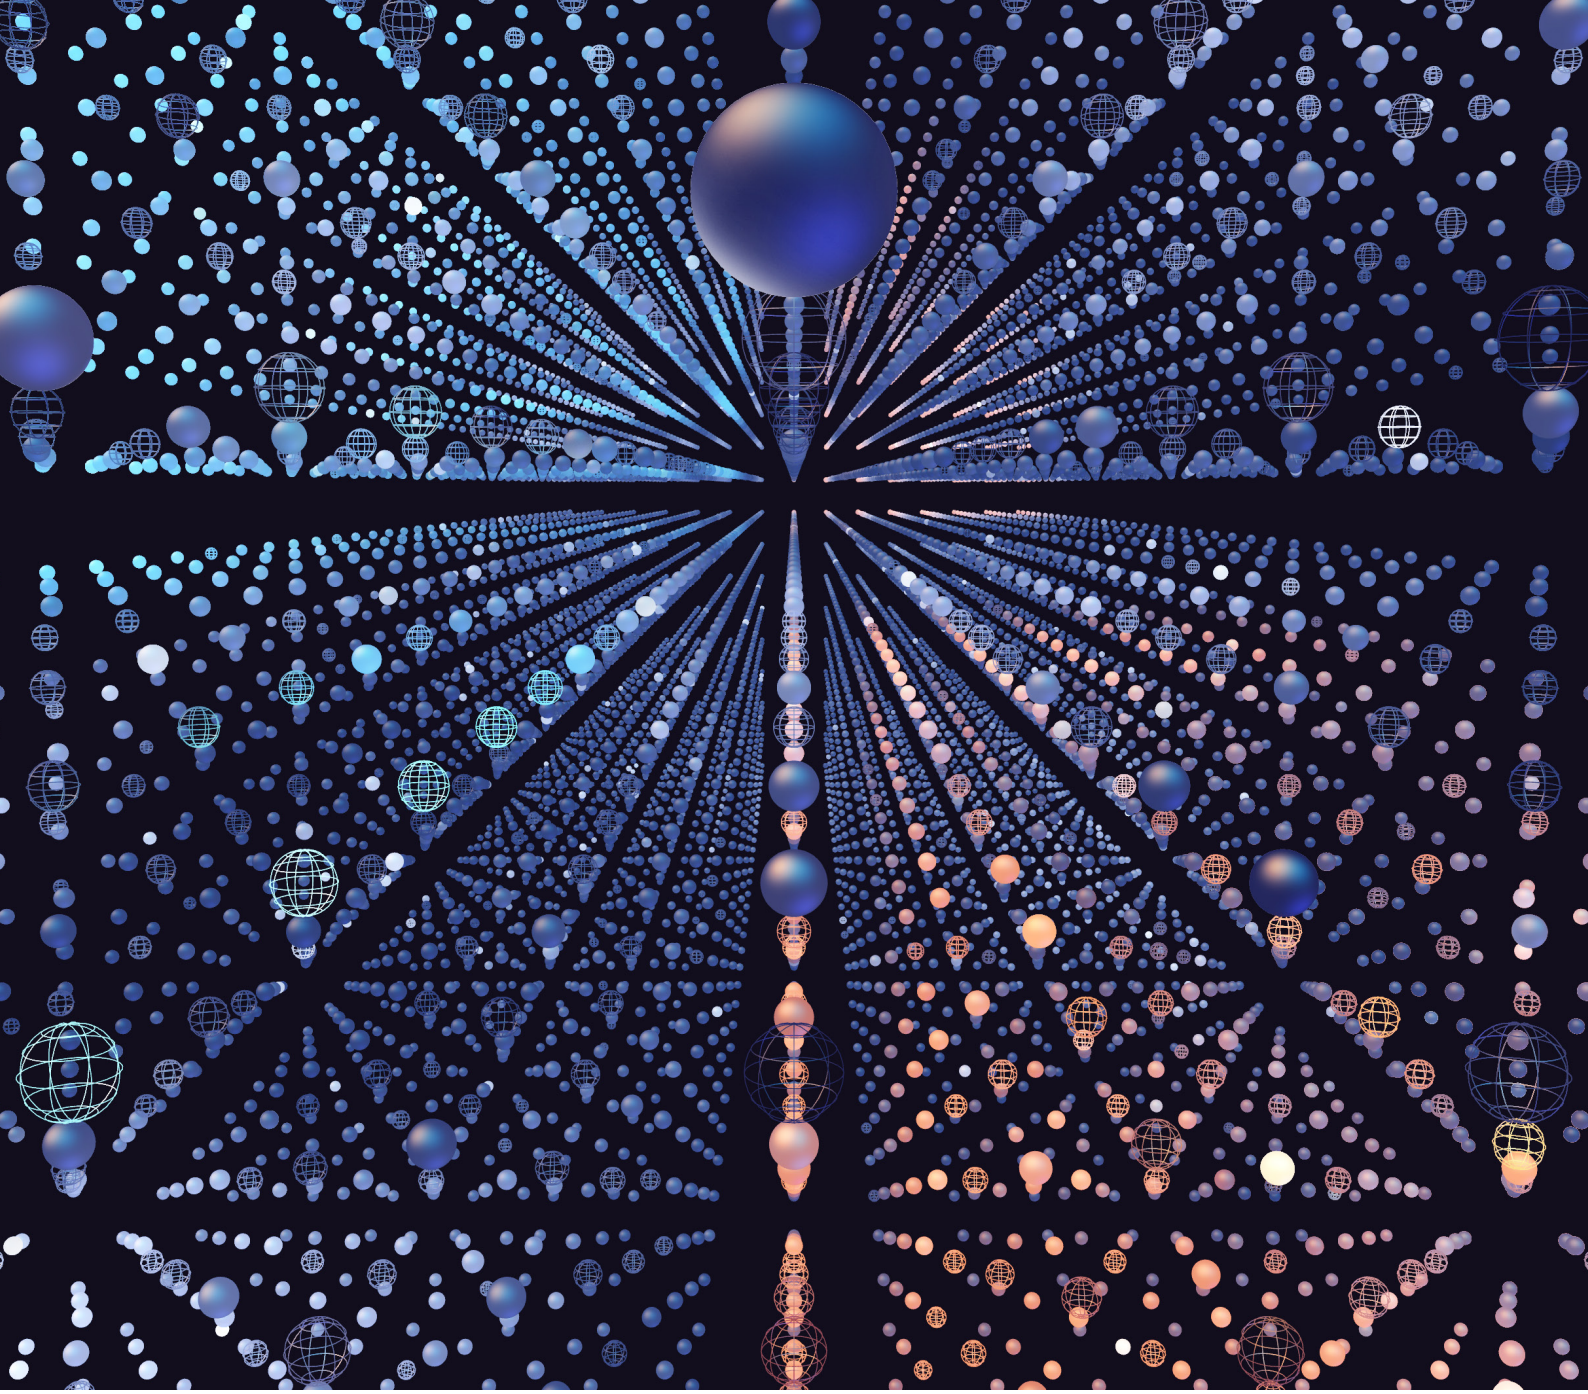

MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com