HARVARD BUSINESS REVIEW

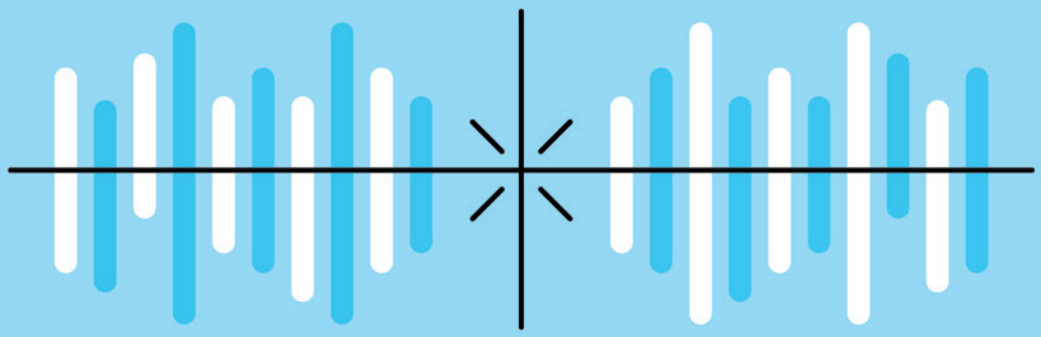

Solving Agentic AI's Data Infrastructure and Telemetry Needs

HBR Analytic Services **PULSE SURVEY**

Sponsored by

Cribl

Telemetry: The Missing Link in Agentic AI

Agentic AI is moving from experimental to operational inside most businesses. Companies aren't just autocompleting emails or summarizing documents anymore. They're asking software to decide what to do next, to stitch together workflows across systems, and to act at machine speed. Such uses of technology represent an epochal shift in how work gets done. It is also a shift in what we demand from our data architecture and, especially, our telemetry.

Cribl sponsored this Harvard Business Review Analytic Services report because we see the same pattern over and over: forward-looking executives have big ambitions for agentic AI, but they're trying to run those agents on top of fragmented, expensive, and opaque telemetry. Most businesses don't have an artificial intelligence (AI) problem in the abstract; they have a telemetry and platform problem in practice. When agents start reading from and writing to the same legacy systems teams already struggle to manage, the volume, velocity, and sensitivity of telemetry not only increase but also change the economics, the risk profile, and the pace of innovation.

Telemetry is not exhaust. It's the real-time narrative of how systems, users, and agents behave, and it's the foundation for the next generation of products and internal tools. An AI-ready data engine must be able to collect, process, transform, analyze, and store telemetry in ways that make it easy to build on top of it. That remit means enabling teams to create their own agents and custom applications, use a built-in development environment instead of rigid on-rails systems that force users into a specific workflow, and rely on federated search as the connective tissue across tools. In that model, organizations can ask and answer questions of their telemetry no matter where it resides and turn those answers directly into new capabilities.

Getting there requires an architectural choice, not just another tool purchase. An AI-ready telemetry foundation has to be open and interoperable so users are not locked into today's vendors. It also must be cost-efficient enough to allow aggressive experimentation without letting telemetry bills and data movement dictate the AI roadmap. The foundation has to provide choice, control, and flexibility of how data is managed and analyzed.

Cribl, the AI Platform for Telemetry, was built on those principles. We believe the organizations that will win in the agentic era are the ones that treat telemetry as a strategic asset with a data engine that their teams and partners can build on, not just plug into. We're pleased to sponsor this research in association with Harvard Business Review Analytic Services, and we hope it helps forward-looking executives connect their AI ambitions to the telemetry, governance, and platform capabilities required to make those ambitions real.

Abby Strong
Chief Market Officer &
Chief Customer Officer,
Cribl

Solving Agentic AI's Data Infrastructure and Telemetry Needs

Organizations are increasingly adopting agentic AI, which can change the nature of what artificial intelligence (AI) can do for businesses. Rather than merely producing summaries and writing text or code, AI agents can plan and execute multistep workflows and learn from context and experience to improve their future performance. That capability can unlock a world of possibilities for rethinking processes and creating business value.

BUT WHILE EXECUTIVES recognize the potential of agentic AI, many don't yet have a strategy for it. "AI and agents are one of the most defining innovations that we've had for 30 years," says David Crawford, chairman of the Technology, Media & Telecommunications practice at Boston-based consulting firm Bain & Co. "You absolutely need a strategy, but the question is what your goals are for these tools."

For many organizations, that prerequisite seems to be the crux of the issue when advancing on their agentic AI journey. In February 2026, Harvard Business Review Analytic Services surveyed 351 respondents from the *Harvard Business Review* audience, all from organizations considering, piloting, or actively using agentic AI and who are both involved in their organization's agentic AI decisions and familiar with telemetry. Ninety-six percent of respondents say they expect the technology to be very important or somewhat important to their organization's business strategy in two years. Yet only 23% say their organization currently has a strategy in place for agentic AI.

The disconnect may be rooted in the fact that setting strategic goals is only the first step. To get the most from agentic AI, companies also must have a data infrastructure that meets the technical needs of their agents. Such an infrastructure should support continual, real-time access to the data that agents require to act, since slow or fragmented systems will hamper agents' effectiveness and decision making.

Data infrastructure is costly, however, particularly when it comes to agentic AI's needs involving

HIGHLIGHTS

96%

of survey respondents say they expect agentic AI to be very or somewhat important to their organization's business strategy in two years.

80%

agree that their organization expects its data infrastructure to change because of agentic AI.

76%

say their organization's volume of telemetry data has increased as a result of using agentic AI.

Due to rounding, some figures in this report may not add up to 100%.

> **"AI and agents are one of the most defining innovations that we've had for 30 years. You absolutely need a strategy, but the question is what your goals are for these tools."**
>
> David Crawford, chairman of the Technology, Media & Telecommunications practice at Bain & Co.

infrastructure upgrades and telemetry, the automated process of collecting and analyzing data about activity in companies' systems. Telemetry is a crucial factor because firms' volume of telemetry data can increase greatly as agents reason, plan, and execute complex workflows. In the survey, of those respondents from organizations that have agents in production, pilots, or experiments, 47% say their organization's data infrastructure costs are higher than expected after adopting agentic AI.

Then there's the need for good governance to dictate how agents are allowed to act. Agents' speed and autonomy mean that companies must have strict guardrails around what they can access and do—and the ability to stop them when necessary. Because agents can potentially make changes to business records and data sources, these concerns go hand in hand with considering liability and compliance.

Taken together, these factors challenge organizations when advancing their agentic AI aspirations. "Without the right infrastructure, you'll hit a ceiling," says Ryan Kurt, CEO of Cleveland-based strategy and advisory firm The AI Lab. "There is absolutely no way to break through it unless you have the data scaffolding, the governance, and the integrated workflows that you need."

This report explores the technical challenges and financial costs of agentic AI. It covers the hurdles that companies face in adopting the technology, the data infrastructures they need for agents' success, the telemetry issues that agents raise, and what IT and security teams should consider about these factors. It also examines agentic AI's financial costs and the governance concerns that executives should be thinking about.

Agentic AI's Strategic Opportunities

AI tools have been in use for several years, but agentic AI can expand a company's toolkit with agents' ability to autonomously carry out a series of tasks, learn from their environments, and improve their approaches for the future.

That dynamism is what differentiates agentic AI from other technological advances—even traditional AI, according to Bain's Crawford. He underscores agentic AI's catalytic abilities with a comparison to cloud computing, a technology that offered productivity and flexibility gains but didn't fundamentally alter how work gets done. "Everyone in the enterprise kept doing their job pretty much the same way with cloud," he says. "Agentic AI is not like that." Agents offer possibilities to redesign operations from the ground up, but Crawford cautions that if executives shy away from doing that hard work, they'll merely be "paving the goat trails"—incrementally improving existing workflows instead of inventing better ones.

While it's early days for agentic AI, the organizations moving quickly on the technology are seeing promising results. AI Lab's Kurt says that most of his clients are at the beginning of the AI maturity curve, but a small number have broken away from the pack by embracing agentic tools and deploying swarms of agents to execute tasks. "Those organizations are creating an environment where a small team of people can now accomplish the work of a very large team at much higher levels of quality, consistency, and communication," he explains.

That spirit of exploration is widespread. The survey shows that agentic AI currently has more potential areas of usage than use cases in production, a divide poised to narrow in the near future. For all business areas listed in the survey, among the group surveyed from the *Harvard Business Review* audience, the percentage of respondents who say their organization is considering using the technology is greater than the percentage who say it's currently in use. The difference is clearly seen, for example, in IT operations (31% using, 52% considering), customer service (29%, 51%), engineering (27%, 39%), product development (24%, 46%), sales/marketing (23%, 53%), and cybersecurity (21%, 44%). At the same time, for different areas of the business, between 9% and 24% say they don't know whether their organization is using or considering using agentic AI, suggesting there is some uncertainty

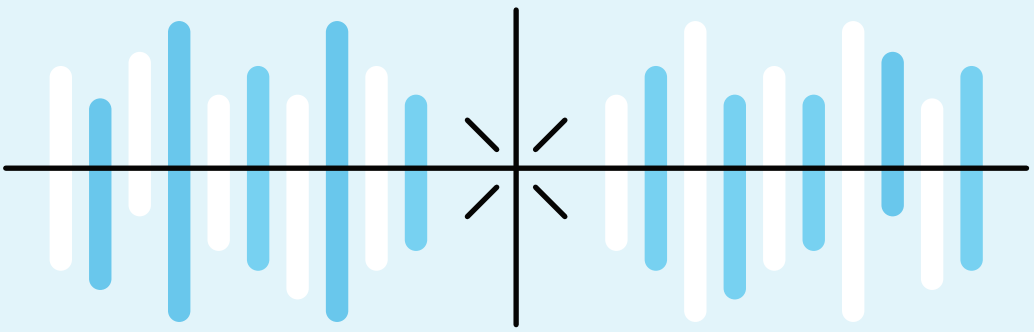

"Without the right infrastructure, you'll hit a ceiling. There is absolutely no way to break through it unless you have the data scaffolding, the governance, and the integrated workflows that you need."

Ryan Kurt
CEO
The AI Lab

> **"It's one thing to deploy these technologies, but it's another thing to drive adoption. There's a lot of discussion around what we need to do to drive the right outcomes."**
>
> Stephen Elliot, group vice president of I&O (infrastructure and operations), cloud operations, and DevOps, IDC

about—and perhaps misalignment of—plans to deploy the technology.

Although exact usage numbers for the fast-moving technology may be hard to pin down, companies will benefit from having a strategy in place to guide their efforts with it.

Of course, without an agentic AI strategy, all the experimentation can become empty effort. So, any such strategy needs to be multilayered, according to Stephen Elliot, group vice president of I&O (infrastructure and operations), cloud operations, and DevOps at Boston-based market research firm IDC. He points out that while the board and the C-suite may understand the AI strategy, it's important that other stakeholder groups do too, from employees who own agentic projects to function managers who identify processes for redesign to AI vendors who support the company's goals.

The other piece of the puzzle, Elliot adds, is to focus on the desired business outcomes of agentic AI, as pursuing the technology for its own sake can lead the company in directions that don't add value. "Companies have to be flexible in the timelines for their AI goals but also recognize the outcomes they want to achieve and who the strategic partners are that they need to really drive these successes," he says.

Challenges of Adoption

Far from being a plug-and-play software purchase, agentic AI represents an epochal shift in how data, people, and processes work together. The technology's potential to transform workflows and the creation of business value is multifaceted, so it's no surprise that adopting it is a holistic effort and brings a range of challenges.

Those hurdles start with trust, because if employees don't see that agentic AI consistently produces its intended outcomes, they're likely to resist using it. Yet trust itself is also multifaceted because agents' success depends on several factors, including how well the company has set up its data infrastructure, governance, workflows, and strategy. The more concretely executives define these parameters, the better agents' results will be and the more employees will be able to trust it. "It's one thing to deploy these technologies, but it's another thing to drive adoption," says Elliot. "There's a lot of discussion around what we need to do to drive the right outcomes."

The survey shows that organizations recognize the breadth of agentic AI's hurdles, which span compliance, skill, infrastructure, and process factors. When respondents are asked what significant challenges their organization has experienced in its efforts to adopt agentic AI, five are mentioned by more than half: privacy risks/issues (57%), lack of necessary talent/skills (53%), data architecture isn't prepared (53%), lack of a strong roadmap/adoption plan (53%), and workflows aren't set up for agentic AI use (52%). FIGURE 1

Whether data architectures are ready to support agentic AI is a crucial element for executives to consider, since agents' long-term effectiveness depends on it. Kurt suggests thinking of that effectiveness on a scale from zero to 10. While organizations may achieve results in the lower end of the scale by, for instance, smartly defining use cases, he says they're unlikely to push beyond that level unless they address agentic AI's technical requirements.

Executives should take note because the survey shows that in many organizations, data infrastructures aren't ready for agentic AI. Respondents are nearly evenly split on the question, with 42% agreeing that their organization's data infrastructure is prepared to meet the technical demands of agentic AI and 43% disagreeing. There's consensus, however, that data infrastructures will change because of the technology; 80% say their organization expects it will.

Once companies deploy agentic AI, a further challenge is telemetry. Having observability into agentic AI's actions and decision making is hugely important because agents operate at blindingly fast speeds. Their use can create risks from poor decisions being carried out instantaneously to bad actors or low-quality data sources influencing their planning. Unless companies can carefully monitor what agents are doing, they'll struggle to stay ahead of potential

negative outcomes—and to respond when they happen. "You want to make customers happy and keep business flowing," says Marc Huguenin, chief development architect of SAP Business Technology Platform Fabric at Walldorf, Germany-based enterprise technology provider SAP. "If you don't have good telemetry, then the time to resolve problems is longer because your people don't have good insights into what's happening." And while legacy telemetry tools capture what happened in their systems, they often fail to explain why, and it's the why that AI agents need to know to keep improving.

Technical Gaps in Data Infrastructures

The specifics of preparing a data infrastructure for agentic AI can depend on how advanced a company's tech stack is and what use cases executives are prioritizing. A further consideration is to what degree an organization will build its agentic AI tools or partner with a vendor for them. So, any investment in data infrastructures should be driven by the company's agentic AI strategy and tied to the capabilities that will support its goals.

For these reasons, Bain's Crawford says that agentic AI efforts and infrastructure upgrades should be business-led, not simply determined by IT or any single function. Some executives may think it makes sense for a technical function to take charge, but that mindset can encourage a too-narrow view of agentic AI's enterprise-wide potential, preventing the organization from laying the groundwork to harness its greatest benefits. "Where our clients have designated this technology to be an IT trend and not a business process trend, they've gotten worse results," Crawford says. "It has to be a business-led activity and supported by legal, security, and other relevant functions."

According to the survey, three aspects of data infrastructures are the most in need of upgrades to meet the technical demands of agentic AI: data quality/hygiene (39%), data integration across systems (37%), and governance/ownership of data (34%). The other data aspects in the survey were mentioned by fewer than one-quarter of respondents.

Companies wondering how to assess their data infrastructure's readiness can begin with those three areas. After all, agents' effectiveness will be limited if they don't have access to clean, high-quality data and the ability to work smoothly across the firm's systems. Cleaning data isn't the

FIGURE 1

Agentic AI Adoption Faces a Plethora of Challenges

Privacy risks are the biggest one, but several others aren't far behind

What significant challenges has your organization experienced in its efforts to adopt agentic artificial intelligence (AI)? *Select all that apply.*

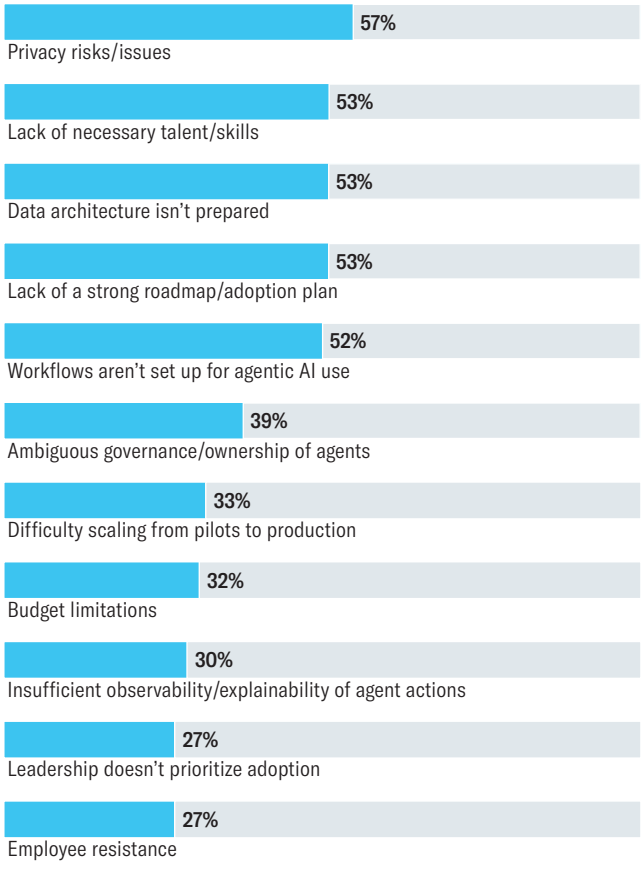

Base: 351 respondents. Not shown: 8% Other, 1% None, 1% Don't know.

Source: Harvard Business Review Analytic Services survey, February 2026

FIGURE 2

Negative Impacts of Unprepared Data Infrastructures

The worst are unclear performance and ROI metrics for agentic AI

What impacts has your organization experienced, if any, as a result of its data infrastructure not being fully ready for agentic AI? *Select all that apply.*

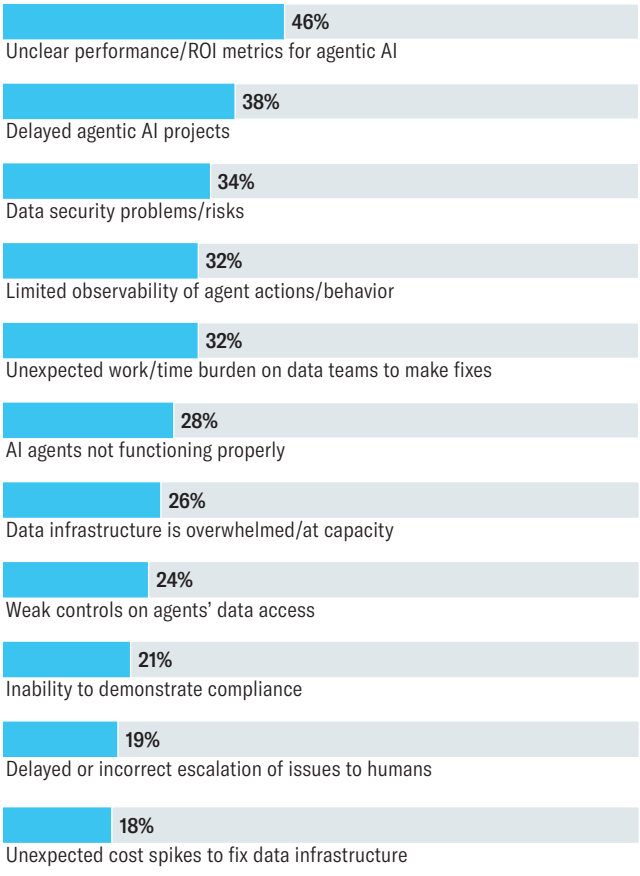

Base: 351 respondents. Not shown: 11% Don't know, 4% None, 3% Other.

Source: Harvard Business Review Analytic Services survey, February 2026

most exciting goal for executives who want to charge ahead with agentic AI, but as Kurt points out, it's a foundational enabler of AI success. "It's hard for a lot of the companies we work with to envision exactly what their business could look like in the future with AI agents," he says, "but everyone understands that the quality of your data is an essential ingredient."

As for the fallout when data infrastructures aren't fully ready for agentic AI, the most-cited impacts in the survey involve uncertain ROI, delays, security risks, and lack of insight into agents' actions. The impacts most mentioned by respondents include unclear performance/ROI metrics for agentic AI (46%), delayed agentic AI projects (38%), data security problems/risks (34%), limited observability of agent actions/behavior (32%), and unexpected work/time burden on data teams to make fixes (32%). FIGURE 2

For companies with data infrastructures that lag behind those of peers, Kurt has some words of encouragement. He's seen that firms with worse tech stacks, perhaps counterintuitively, can turn into infrastructure bellwethers. In more-advanced companies, decision makers often focus on plugging agentic AI into existing systems, which can hinder the work of reimagining processes. But when tech stacks are less advanced, decision makers have more leeway to design workflows from scratch. "There are fewer guardrails and less bias to integrate AI into the way you already do things," he explains. "So instead, people are creating brand-new processes where the goal is to solve a business problem."

Expanding Telemetry Volume

Going hand in hand with reinforcing data infrastructures is companies' need to bolster their telemetry capabilities. Closely monitoring agents' actions is inescapably critical due to the speed they act with and the cross-systems nature of their work. Without good observability, executives won't know what their agents are doing or which data sources and reasoning approaches they're using to do it.

The key point for decision makers to understand is that the more expansive their data infrastructures, the greater their telemetry volume. SAP's Huguenin points out that while telemetry has tended to increase with each evolution of the tech stack, the volume implications for agentic AI are an order of magnitude greater. When companies moved from monolithic architectures to microservices, for example, telemetry volume increased—say, by a factor of 10—because there were many more services to monitor. "Now you're going to use agents to make those microservices smarter, and you'll exponentially increase the amount of telemetry you'll have to process from the agents," Huguenin explains, adding that the monitoring challenge for companies is both clear and urgent.

Huguenin says that while agents likely will increase companies' volume of telemetry data, agents are also what will enable organizations to handle that volume. Previous methods of monitoring, such as data dashboards, lack the scale, speed, and breadth to remain sufficient. "When you [increase by] 10 times or 100 times the amount of data coming in, dashboards aren't a solution anymore," Huguenin explains. "You need to use agents to look for anomalies in the data and then take action."

The survey finds that agentic AI is already increasing the volume of telemetry data for most. Of those respondents from organizations with agentic AI in production, pilots, or experiments who also confirmed telemetry is in place in their organization, a majority (76%) say their organization's volume of telemetry data has increased. A plurality (45%) say telemetry data has increased slightly, while 31% say it's doubled or more. Meanwhile, 24% say their volume of data has not changed, suggesting not all are seeing the data impacts yet. FIGURE 3

As for handling that volume, the survey also shows that the use of agentic AI in telemetry has room to grow. Just over one-quarter of respondents (27%) say their organization is using the technology in telemetry, with 14% saying it's a core part and 13% using it in limited areas. Around half of respondents (51%) say their organization is experimenting with agentic AI in telemetry (23%) or considering doing so (28%).

FIGURE 3

Pumping Up Telemetry Volume

The majority of respondents say their agentic AI efforts have led to an increase

In your estimation, how has your organization's volume of telemetry data changed, if at all, as a result of using agentic AI? *Select one.*

Because of agentic AI, my organization's telemetry data has:

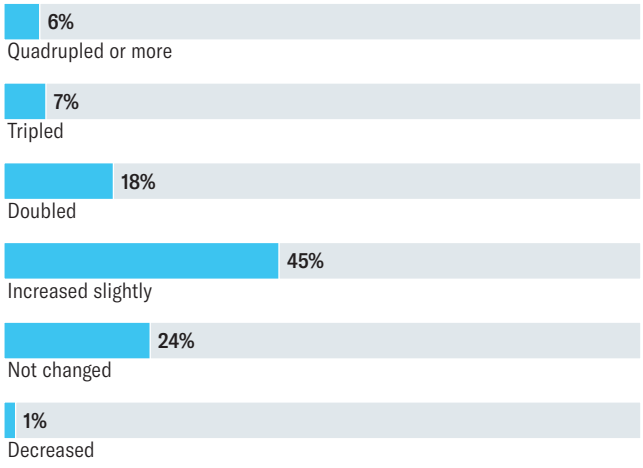

Base: 190 respondents whose organizations have agentic AI in production, pilots, or experiments and who report having telemetry in place at their organization. Not shown: 0% Don't know.

Source: Harvard Business Review Analytic Services survey, February 2026

Telemetry can also help organizations spot new opportunities to either automate a process or flag its inefficiency. Crawford, for instance, suggests letting agents identify the key "nouns" of the company—supply chains, purchase orders, delivery vehicles, and so on—as well as their relationships and interactions with graph analytics and ontologies. Those insights can help decision makers consider how effectively those resources are being used. "You end up having a very strategic asset in the capacity to study what the organization is actually doing and where it's spending its resources," he says. "That opens up new possibilities of creating candidates for automation or generating agents that can drive value for you."

Kurt is similarly enthusiastic about agent-driven telemetry, suggesting two ways to expand its effects. The first way involves mapping how teams work, which many teams aren't great at documenting, he says. The goal is to get visibility into and then articulate a group's standard operating procedures, which can help train agents to handle both common and uncommon scenarios. Such a capability could

> "There's never been a better time to step back and say, 'How should this get done?'"
>
> The AI Lab's Kurt

prove invaluable when responding to security risks, for example. "Your team can sit down and define, 'When this happens, this is how we want to handle it step-by-step,'" he explains. "Then you can give that to an agent and improve its ability to act more autonomously and apply your policies at scale."

The second area Kurt mentions is in searching for process improvements. Because teams often don't have playbooks for their work, articulating their usual methods can also help them consider where better outcomes are possible and how to achieve them, drawing on telemetry data for those insights. Executives shouldn't forget that merely automating the company's tried-and-true procedures is unlikely to help agentic AI create as much business value as possible. "There's never been a better time to step back and say, 'How should this get done?'" Kurt asserts.

Finally, good telemetry will help organizations optimize their agents' performance, according to Huguenin. Development and operations teams have long used telemetry data to identify problems to resolve for customers. Such data can also help them see where agents may be underperforming and dig into the causes of it. "With the power of AI agents and the benefits you get from them, injecting that data back in makes them smarter," he says. "Figuring out a solution for that makes a lot of sense if you want to be best in class."

Reshaping IT and Security

Beyond telemetry, organizations have high hopes for how agentic AI can improve their IT and security functions. Integrating agentic AI with those teams' core work can bring new opportunities in human-AI collaboration, such as spotting system vulnerabilities and improving feedback loops to make processes more robust.

The benefits that organizations seek in IT and security, though, largely come down to simplifying and accelerating work. The survey shows that of those at organizations using or considering using the technology in IT/security, the top two hoped-for benefits are streamlined processes—i.e., fewer and simpler steps—and faster detection of problems, both mentioned by 36% of respondents. Other sought-after benefits include tech teams having more time for high-value work (29%), faster reaction to problems (28%), and less need for manual intervention (22%).

A question that may be top of mind for executives is to what degree agentic AI will replace technical employees versus taking over specific tasks for them. According to Huguenin, the latter is the more common approach—that is, instead of supplanting employees, agentic AI should be used to amplify their abilities. "I don't think there's really any intention to replace humans with agents," he says. "We use agents to make people smarter and faster at what they do."

He describes this form of collaboration as a bidirectional data flywheel. Essentially, employees use apps in their work, generating data that is fed into agents. The agents use that data for their own tasks, creating more data that is fed back into the apps. All the while, humans are kept in the loop to monitor processes and make decisions where needed. "That cycle is well-defined at SAP and is already established with all our agents," Huguenin explains. The data flywheel helps improve the performance and effectiveness of both agents and apps, generating higher-quality data throughout the cycle.

Another opportunity in IT and security is spotting system vulnerabilities before bad actors can exploit them. It's essential to patch these areas as soon as possible, since the longer they stay open, the greater the company's exposure is. Agents fit that need well because they can monitor and respond to danger orders of magnitude faster than humans can. Pair that speed with agents' reasoning and learning abilities, and organizations' security response could become more potent than ever. "Where a lot of this will start to head is not only looking for risks and responding to incidents as they happen," Kurt says, "but starting to predict where there might be a vulnerability that hasn't even opened up yet."

At the same time, agents can become security risks themselves because of their access to systems and the potential

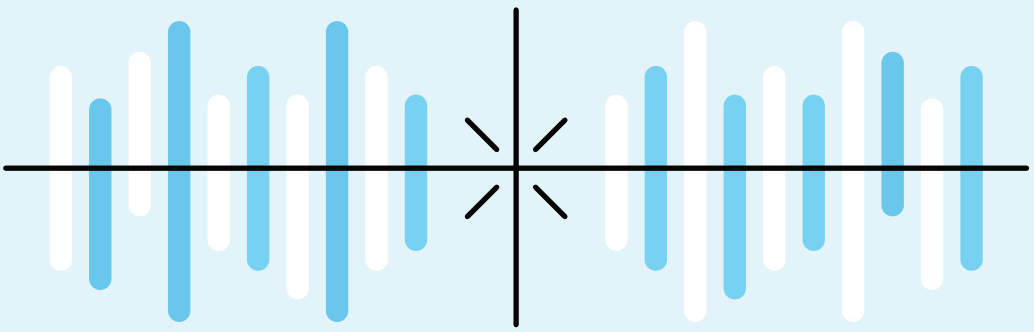

"We use agents to make people smarter and faster at what they do."

Marc Huguenin
Chief development architect
SAP

difficulty of understanding their actions. IT and security teams must know how to spot unusual behavior, drawing on telemetry data and other sources of insight to determine whether agents behave consistently and produce consistent results over time. Here again, agents' speed is an enormous factor, since an agent that creates a security risk will move infinitely faster than a human would. "What's the model or security paradigm for detecting those kinds of problems?" Huguenin says. "The security concepts for agents are being heavily thought out at SAP."

Prioritizing Governance

To ensure agentic AI is working safely and compliantly, executives will have to put in place the proper governance for it. Doing so involves making strategic decisions about its use cases, access to systems, legal liability, and decision making, as well as considering how AI vendors weigh these factors.

While governance is vital, Crawford says it's difficult to think about abstractly, so the conversation must be tied to how a company wants to deploy agentic AI. "IT and legal teams have an important role in monitoring and setting policy, but it's hard for them to do that without knowing what teams are going to do with it," he explains.

If a company implemented a tool that provided content and data to its salespeople to save them time, for instance, the legal team might not object, given the impact's straightforward nature. If a tool monitored salespeople's communications to suggest potential customers for them to contact, more complicated issues could arise quickly. The lesson thus is that governance conversations should be moved from the abstract to the concrete as soon as possible.

More broadly, Crawford says it's paramount to have governance around which information sources an agent can access, how agents communicate both within the company and outside it, and what kinds of actions agents can take. Instituting the right controls in these areas is a critical aspect of proactively addressing cyber risks. "These things are nontrivial, particularly in a world where there's legal liability involved," he explains. "There are massive amounts of hackers around the globe trying to take advantage of whatever you do and drive their own agents accordingly."

Also important is to determine what kinds of decisions agents can make and when humans must be in the loop to advise or make the decision themselves. Elliot points out that while executives may leave smaller choices up to agents, larger or more nuanced choices should invite caution. As an example, he points to the difference between resetting login credentials and selling a stock in an investment bank's portfolio, which may involve strict compliance regulations. "Passwords are very important, but resets are a pretty straightforward process, so you could probably automate that and not have any humans engaged in it," he says. "But if an investment firm is selling a stock, how many people would leave that decision up to an agent?"

Elliot emphasizes that executives should also learn the governance practices of AI vendors they work with to be sure the companies' security norms align. After all, different agents will have access to different types of information and be governed by different companies' guardrails. "That's going to drive a lot of complexity in terms of how you govern, how you orchestrate—and how you shut agents down when something isn't going right," he says.

An additional concern is establishing the right evaluation frameworks for agentic AI. As Kurt explains, organizations running agents in high-risk business areas should regularly assess their performance against benchmarks. Many companies monitor the performance of physical assets, such as factory machines. Agents require similar evaluation so that decision makers know the technology is continuing to produce its expected level of quality. "Think of it like an airplane engine," Kurt explains. "The more important that engine is for your company, the more you need to monitor it on a frequent basis so you can know when it starts to behave a little differently and intervene when necessary."

A final governance caution has to do with capturing employees' expertise to train agents. While Kurt sees the

> **"Our clients who have initiated big horizontal data reinventions have been pretty quickly frustrated with the amount they've spent and the absence of new capabilities that they've enabled."**
>
> Bain & Co.'s Crawford

benefits of mapping processes and specialists' knowledge to make agents smarter, he thinks it could fundamentally change the traditional employer-employee contract. "You're essentially asking employees to partake in a process of building a digital twin of how your company operates," he says.

Companies that take this route are no longer simply leasing employees' know-how in exchange for a salary and benefits; they effectively are using that know-how to create permanent assets under the company's control. The full implications of that shift may not be known for years. "Once you load information from people's brains onto a server, you're turning employees' knowledge into an asset that the company owns," he continues. "That brings up a lot of ethical and legal questions that I think are going to become very pertinent over time." Kurt suggests executives have clear conversations with employees about any such mapping and how it affects the workforce.

The Financial Side of Agents

As organizations continue their agentic AI journeys, it's natural to question the implementation costs of the technology. Agentic expenses can quickly grow out of hand if decision makers aren't careful, and executives should be wary of underestimating them—perhaps drastically.

The best starting point is to strategically tie the company's financial investments for agents to the business goals it wants to achieve with them. What Crawford has seen is that organizations moving ahead without a roadmap for realizing business value often fail to capture enough ROI. "Our clients who have initiated big horizontal data reinventions have been pretty quickly frustrated with the amount they've spent and the absence of new capabilities that they've enabled," he says.

A better approach is to incrementally invest toward desired outcomes, such as by consolidating databases or sunsetting legacy systems. That approach won't lead to immediately building all the capabilities that companies will need, but it can help them avoid overspending while pinning down their agentic AI goals.

Avoiding excessive expenditures is important because the survey shows that the data infrastructure costs of agentic AI often exceed expectations. Nearly half of respondents (47%) from organizations that have agents in production, pilots, or experiments say their organization's data infrastructure costs are higher than expected after adopting agentic AI, with 22% saying theirs are about equal to expectations. Only 7% say their costs are lower than expected, and 24% say the impact is unclear and cite "don't know." FIGURE 4

However, even if all organizations haven't felt an increase yet, four in five see it coming: 82% of respondents agree with the statement "My organization expects there will be a financial cost to meet the data infrastructure demands of agentic AI." Additionally, half of respondents say expected infrastructure costs are included as part of their organization's total cost of ownership (TCO) for agentic AI, with 11% saying they aren't included. Another 29% say their organization has not calculated a TCO for agentic AI. These findings suggest that some executives underestimate the TCO of agentic AI by not considering back-end upgrades such as the observability tools, AI data platforms, and low-latency network architectures required to support it.

For Huguenin, managing and optimizing the TCO for SAP's customers is a continual goal. The financial factors include having the right talent, investing in the right platform and efficiency capabilities, and measuring ROI the right way, which can include cycle times, reliability, and user productivity. He's seen that agentic AI can be an expensive endeavor, but he emphasizes that those costs should be viewed as an investment. "You're making your apps a lot smarter with AI agents," he explains. "The key point is what you get out of it, how you optimize the ROI."

Another component of agentic AI costs is whether companies partner with vendors. Key considerations, according to Elliot, include the AI models vendors use and whether those vendors own all their own infrastructure,

FIGURE 4

The Costs of Adopting Agentic AI

Nearly half of respondents have seen their organization's data infrastructure costs rise

Compared to your organization's expectations, how has adopting agentic AI impacted its data infrastructure costs, if at all? *Select one.*

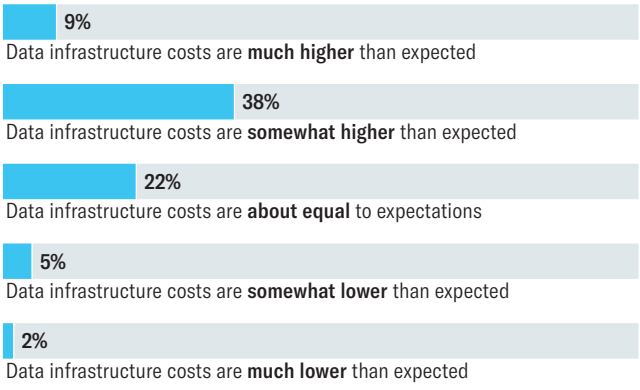

Base: 218 respondents whose organizations have agentic AI in production, pilots, or experiments. Not shown: 24% Don't know.

Source: Harvard Business Review Analytic Services survey, February 2026

which may influence the TCO. In addition, he suggests thinking about which use cases a company might want to keep in-house and which ones could be done through a vendor's systems, since keeping any in-house—such as those involving sensitive information or proprietary data—will require a data infrastructure that can handle them. "You have to think about storage, compute, power, energy, security parameters, data parameters—you layer all these costs on from an AI infrastructure perspective," Elliot explains.

Conclusion

Agentic AI is poised to transform how companies operate, giving them a powerful tool that can autonomously complete multistep workflows and learn over time. Still, all the tasks necessary to position agentic AI to perform—developing adequate data infrastructures, providing access to data sources, and conjuring thoughtful governance guardrails—should be driven by a cohesive strategy that defines how agents will create value for the company.

To capture that value, organizations must have data infrastructures that support agentic AI's technical needs, which is likely to require upgrading legacy systems, whose limits agents will quickly overwhelm. Agents also need access to high-quality data sources and insightful governance to dictate how they can act, determine which decisions they're empowered to make, and evaluate their performance over time. Telemetry, too, should be a priority, since executives must have insight into what their agents are doing and know when adjusting or halting their use is called for.

Executives' to-do lists for adopting agentic AI may feel daunting, but now is the time to get started. Companies are increasingly experimenting with and using the technology, and as their facility with it grows, organizations that delay their investments will find themselves lagging behind agent-powered competitors.

"Sometimes it's good wisdom with a hot new technology to wait out the first wave of innovation, but with agentic AI I think there's a real risk around that," says Crawford. "You're putting yourself at a significant risk of being disrupted by someone who's done more business redesign work with it."

METHODOLOGY AND PARTICIPANT PROFILE

Harvard Business Review Analytic Services surveyed 351 members of the *Harvard Business Review* audience via an online survey fielded in February 2026. Respondents qualified to complete the survey if their organization is considering, piloting, or actively using agentic AI; they are involved in their organization's agentic AI decisions (including decisions around data); and they're familiar with telemetry.

ORGANIZATION SIZE

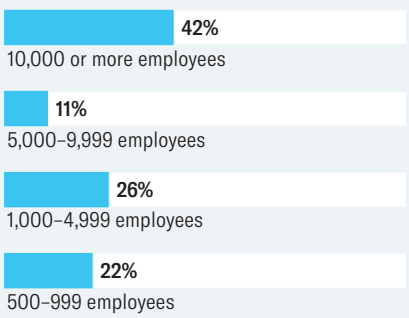

SENIORITY

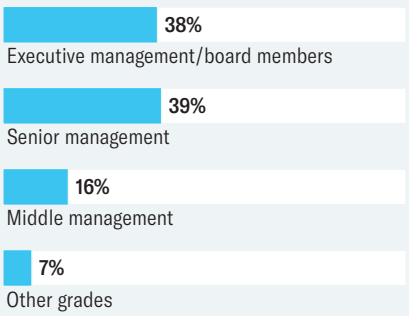

INDUSTRIES

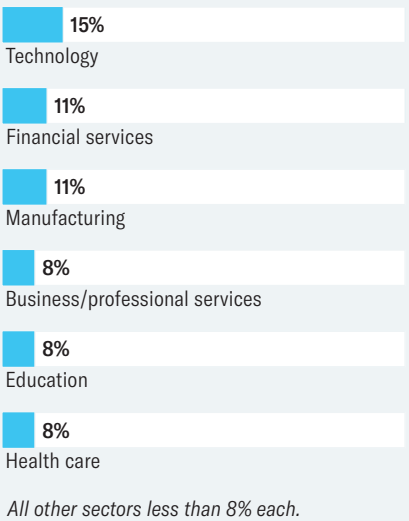

All other sectors less than 8% each.

JOB FUNCTIONS

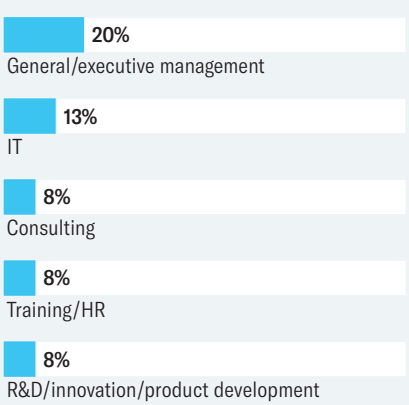

All other functions less than 8% each.

REGIONS

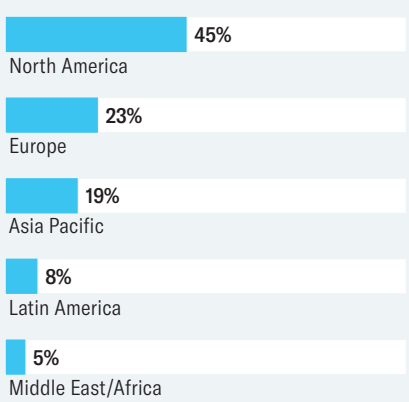

HARVARD BUSINESS REVIEW

VISIT US ONLINE

hbr.org/hbr-analytic-services

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.

CRE77000426