

Intelligent Operations

*How Agentic AI Is Aiming
to Reshape IT*

HBR
Analytic
Services
**WHITE
PAPER**

Sponsored by



Reimagining Cloud Operations in the Era of AI

Cloud operations are entering a new era. It's no longer about automation alone; it's about systems that are intelligent, adaptive and agentic. Intelligent operations turn data into knowledge to optimise resources, adaptive systems anticipate and respond dynamically to change and agentic AI capability enables the autonomous execution of your intent across environments. Together, these cornerstones redefine how organisations manage complexity and deliver agility, pioneering a new way to run operations in the cloud.

IT professionals are facing surging demands and complex environments that call for intelligent systems that can think and act autonomously. Agentic AI cloud operations have become the core foundation for a new operational paradigm – one that shifts from reactive automation to proactive resolution.

Leveraging advanced technologies such as artificial intelligence (AI), machine learning and data analytics, IT professionals are transforming IT through intelligent operations. This approach enables organisations to proactively monitor, optimise and secure their cloud infrastructure in real time, leading to improved efficiency and reduced operational risks. By automating routine tasks, detecting anomalies and providing predictive insights, intelligent operations empower

IT teams to respond swiftly to issues, enhance system performance and ensure compliance across diverse environments. Additionally, this methodology facilitates capacity planning, reduces downtime and streamlines incident response, further enhancing the agility and reliability of IT services.

But intelligence alone isn't enough. IT ecosystems must be adaptive, responding to evolving workloads, security threats and compliance requirements. Adaptive operations integrate federated governance and extensibility, ensuring systems scale seamlessly while maintaining precision and resilience. This adaptability empowers IT professionals to react faster, mitigate risks and sustain performance in volatile landscapes.

The future of IT operations is agentic, where autonomous agents evolve from mere tools into trusted assistants. Unlike static automation, agents embody autonomy with accountability, preserving human oversight while managing routine tasks such as reducing mean time to recovery, optimising costs and enabling real-time remediation across ITOps, DevOps, SecOps and FinOps. To ensure these agents operate safely and effectively, IT professionals must lead change management, redefine accountability and cultivate trust by deploying



Annie Pearl

Corporate Vice President,
Azure Experiences & Ecosystem
Microsoft

infrastructure that enables oversight and control of each agent's actions.

We are proud to sponsor this report, reaffirming our commitment to advancing research that shapes the future of cloud operations in the AI era. As agentic capabilities become integral tools for organisations, we hope the insights shared here guide your journey toward a platform that amplifies the benefits of agentic cloud operations.

The future of cloud operations isn't just automated; it's intelligent, adaptive and agentic, and Microsoft Azure is here to support your success.

Microsoft sponsored this study; however, the information, opinions and conclusions presented are those of the authors and do not represent or imply Microsoft's endorsement or position.

Intelligent Operations

How Agentic AI Is Aiming to Reshape IT

Current IT operations rely largely on reactive monitoring and manual intervention to address system issues. Agentic AI is aiming to redefine IT management through artificial intelligence (AI)-driven agents that autonomously observe, reason and act to resolve issues and optimise systems within defined guardrails and human oversight. This approach enables IT teams to not only anticipate problems, but also resolve them autonomously and more quickly across increasingly complex infrastructure.

“Using agentic AI, we can get a step ahead,” says Sanjeev Mohan, a former Gartner Inc. research vice president who now heads his own advisory firm, SanjMo, in San Francisco. “For example, agents can throw hackers off by constantly looking for loopholes and fixing them before bad actors can exploit the vulnerabilities. They are more proactive than other approaches.”

Cloud operations – including ITOps, DevOps, SecOps, DataOps and FinOps – are driving the adoption of agentic AI in enterprises. A June 2025 study by the consultancy Information Services Group (ISG) in Stamford, Conn., found that IT operations represent 52% of agentic AI use cases, significantly ahead of marketing and sales (12%) and finance (10%). **FIGURE 1**

Companies leveraging agentic AI in IT operations are already seeing measurable benefits. One bank cut the work required to modernise legacy IT systems by more than half using agents, according to ‘Seizing the Agentic AI Advantage’, a June 2025 report by New York-based McKinsey & Co. Other firms cited by the consultancy report smoother IT workflows, faster software development cycles and improved collaboration between IT and business teams as AI agents are integrated into core IT processes.

While such results demonstrate agentic AI’s potential, achieving these benefits demands cultural change, updated employee skill sets and governance frameworks that clarify accountability and human oversight. The challenge is that agentic AI spans

HIGHLIGHTS

Cloud operations – including ITOps, DevOps, SecOps, DataOps and FinOps – are driving the adoption of agentic AI in enterprises.

Effective agentic IT operations depend as much on organisational readiness as on technology.

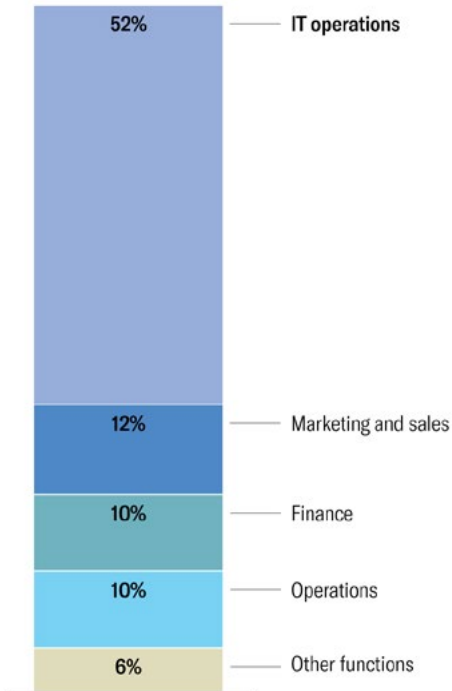
Agentic IT operations are expected to evolve beyond reactive and predictive functions toward automatically improving system performance across multiple cloud and on-premises platforms.

FIGURE 1

IT Leads in Agentic AI Adoption

More than half the use cases involve IT operations

Agentic AI use cases targeting function-specific processes within enterprises:



Source: Information Services Group, June 2025

a spectrum, from simple automation to complex reasoning systems and each agentic AI attribute requires different implementation approaches.

“Organisations are successfully deploying basic automation agents for straightforward tasks, but intelligent agents capable of reasoning and making complex decisions introduce greater implementation challenges,” says John Thompson, senior vice president of generative AI (gen AI) programs at The Hackett Group, a consultancy in Miami.

This report will explore how IT organisations are adopting agentic AI operations today, what outcomes they’re seeing, what barriers they face and how this shift is redefining operational excellence. It will provide a roadmap for unlocking the next wave of performance and innovation through intelligent, autonomous IT operations across public, hybrid and multicloud environments while addressing industry hesitancy about autonomous agents in production.

Companies leveraging agentic AI in IT operations are already seeing measurable benefits. One bank cut the work required to modernise legacy IT systems by more than half using agents, according to ‘Seizing the Agentic AI Advantage’, a June 2025 report by New York-based McKinsey & Co.

Agentic AI Requires Digital Maturity

David Linthicum, former chief cloud strategy officer at London-based Deloitte Touche Tohmatsu Ltd. and now an independent consultant, emphasises that achieving digital maturity is essential when considering the use of agents.

“Companies often lack a clear understanding of how this technology works. Are their data pipelines established? Is the security infrastructure adequate?” Linthicum asks. “The biggest problem is organisations think AI will fix everything, so they deploy it over insufficient infrastructure, which frequently increases complexity rather than reducing it.”

Sahil Sanghvi, a vice president of AI engineering at Booz Allen Hamilton Inc., a consultancy in McLean, Va., reinforces this infrastructure focus. “Successful deployment of agentic AI requires a company to provision access to the right tools and the right data and ensure that the data is being processed and made available in a way that an agent can actually leverage it,” he says.

This infrastructure concern is shaping how companies approach implementation. Rather than risk damaging systems they’re not sure can handle AI properly, companies are beginning with AI that watches and alerts rather than AI that actually changes system settings. A survey of 406 IT leaders and decision makers by 451 Research, a part of S&P Global Market Intelligence in New York, found that

74% of companies use gen AI in IT operations. Most of this technology focuses on observational tasks such as monitoring (50%) and security scanning (49%) rather than core infrastructure management like configuration (35%) and provisioning (23%). **FIGURE 2**

This cautious adoption pattern reflects lessons learned from previous technology cycles. Mark Settle, former chief information officer at Okta Inc., a security company in San Francisco, and author of the book *Truth from the Trenches*, draws parallels to earlier hype cycles. “It’s just like history all over again,” he explains. “When machine learning first came out, companies hired a bunch of really high-priced data scientists and put them in R&D groups. Nobody really understood what they were doing. They prototyped extensively, and very little came out of this approach and was put into production.”

More mature capabilities – such as a data management strategy, knowledge infrastructure, architecture and orchestration and security and governance – enable companies to move beyond basic automation. When these foundational qualities are present, agentic systems can deliver sophisticated predictive operations by recognising patterns across vast data sets, connecting seemingly unrelated anomalies that humans might never correlate.

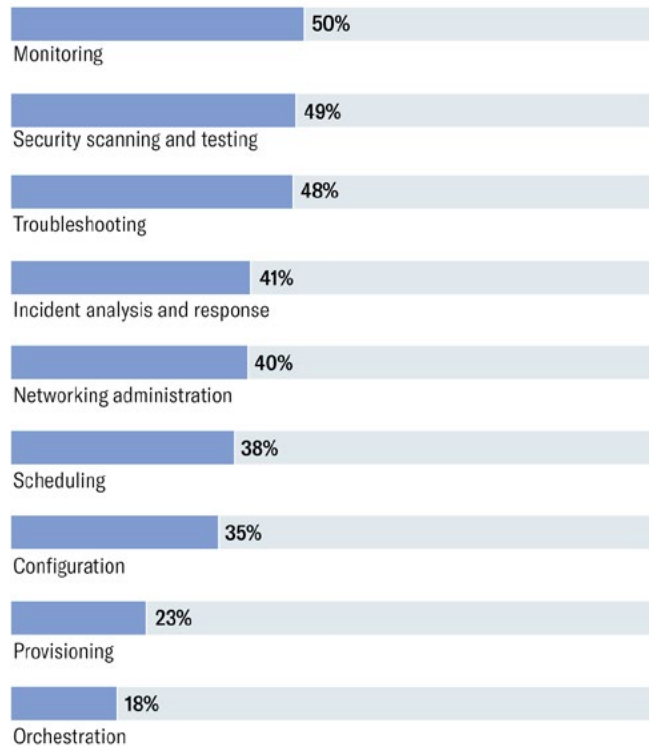
“The agents can determine unusual CPU [central processing unit] spikes, elevated memory usage and increased network traffic across multiple systems,” Settle says. “The real opportunity lies in shifting to predictive capabilities where agents remember prior concerns and link them to new incidents, identifying developing problems before they escalate.”

Predictive intelligence and autonomous capability provide many opportunities to optimise IT operations. Agents can schedule software updates and system optimisations during low-impact periods by analysing historical usage patterns and coordinating complex activities across systems without human intervention. As Settle notes, “System load doesn’t really increase until the United States wakes up, so why not schedule most maintenance between midnight and 5 a.m. eastern time?”

FIGURE 2

Top IT Operations Tasks Using Generative AI Monitoring leads the way, followed closely by security scanning and troubleshooting

IT operations tasks using gen AI (multiple responses allowed):



Source: 451 Research survey, September 2025

Agents could also predict a surge in compute demand, proactively make recommendations and rightsize the virtual machines to accommodate the growth. Another area where pattern recognition delivers immediate operational value is predictive maintenance. The Hackett Group’s Thompson describes implementations where agents monitor server hardware characteristics – such as how hot the server gets, how often it’s turned on and off and how long it’s been running – to anticipate component failures.

“We know that the server has been in its chassis for two years, it has experienced extreme temperatures four times, it’s been shut off many times,” he says. “Agents can compare that to historical data and determine drives in this server are going to fail within six months.”

Agentic IT operations become even more sophisticated when multiple specialised agents coordinate within workflows, enabling complex problem-solving beyond a single agent’s capabilities. SanjMo’s Mohan illustrates a

“ We know that the server has been in its chassis for two years, it has experienced extreme temperatures four times, it’s been shut off many times. Agents can compare that to historical data and determine drives in this server are going to fail within six months. ”

John Thompson, senior vice president of gen AI programs, The Hackett Group

common scenario: a password reset request. A monitoring agent logs the request in the company’s ticketing system, and then an analysis agent determines the action needed, such as whether a password should be changed or an account deactivated.

The system then triggers appropriate specialised agents, which may open tickets, enforce security protocols or ensure regulatory compliance. “A compliance agent can automatically apply GDPR [the European Union’s General Data Protection Regulation] when a user has an account in Germany,” Mohan says. “This orchestrated approach enables multiple agents to work efficiently while ensuring regulatory compliance. Automating mundane operations can improve productivity of the IT teams and allow them to focus on more strategic high-value tasks.”

Setting Boundaries for Agents

Autonomy in agentic AI raises critical questions about control, accountability and risk management. Instead of creating new governance frameworks, experts say many organisations adapt existing policies and procedures for agentic AI. Agents should be treated as virtual assistants with defined roles, rules and expectations, but without attributing to the technology human emotions, consciousness and intention.

“Don’t anthropomorphise agents,” says Thompson. “But I do encourage people to consider agents from a functional perspective. It gets people creeped out, but you can understand agents’ roles if you think of them as synthetic workers.”

Enterprises can institute this governance through natural language understanding, allowing agents to interpret corporate codes of conduct, spending limits and operational guidelines without needing complex rule coding. This approach enables agents to implicitly understand directives such as spending restrictions or ethical guidelines.

Implementation requires careful risk management, particularly during early deployments. Independent

consultant Linthicum emphasises limiting potential damage by restricting agents’ access to critical systems and taking a graduated approach to permissions. He recommends starting with read-only access, where agents can view data and system settings, but cannot modify, delete or create files, security permissions or other settings.

“Put guardrails around who’s accessing the database, who’s accessing the files, who’s accessing the system configuration, because they can really screw things up if they go off-kilter,” Linthicum says. “Capabilities can then be expanded gradually as confidence builds and systems prove reliable, much like onboarding a new employee with high-level system access.”

The degree of oversight varies by use case and regulatory requirements. According to the ISG study, only 25% of current agentic AI systems allow agents to operate fully independently, while 45% function as advisers supporting human decision makers. “A bank using agentic AI for credit and lending decisions will need higher explainability and transparency because of regulatory obligations,” Thompson says.

Practitioners are learning these lessons through experiences that sometimes prove costly. Thompson recounts a cautionary tale where two employees building agentic AI technology inadvertently ran up cloud computing bills of USD 300,000 each within a couple of days. “There was no nefarious intent,” he explains. “They were just excited about what they were doing and their agent ran wild.”

While such incidents were more common in initial implementations, they’ve driven organisations to establish stronger spending controls from the start. “You say to the agent, ‘You only have the authority to spend USD 10 or USD 100 or USD 1,000. Anytime you want to spend more than that, you have to talk to me as a person,’” Thompson advises. This instruction means once the agent reaches its spending limit, it would need to cease its task and send a notification requesting human approval before continuing. This human-in-the-loop approach prevents runaway costs while allowing routine tasks to proceed autonomously.

Measuring Success and Demonstrating Value

Establishing clear metrics for agentic IT operations is crucial for justifying investments and scaling projects, although traditional ROI calculations must be adapted for autonomous systems that continuously learn and improve.

This measurement challenge reflects deeper issues with how organisations evaluate preventive technologies. Traditional IT business cases rely on demonstrable cost savings or revenue increases, but predictive systems can deliver value through problems that never occur.

“If you were to justify a purchase of a tool for its predictive capabilities, the ROI would be like, ‘We avoided this many tickets’, or ‘We never had these problems,’” Settle says. “Stakeholders can have a hard time buying that kind of risk reduction.”

Human psychology related to AI decision making compounds the challenge. Settle observes a dangerous trust bias: “If there were two of us in a conference room and you said this is what we should go do, we’d all jump on you and say, ‘Wait a minute; we got 20 questions before we’re going to do what you think is right,’” he notes. “If AI says here’s the 20 steps, people just say, ‘Okay, what’s step two? Let’s get rolling.’”

Organisations are developing comprehensive measurement frameworks to address this challenge. Experts advocate a three-tiered approach for clarity, technical performance, business impact and organisational adoption.

The first tier consists of basic metrics, including agent uptime and accuracy, that measure technical performance. The second tier involves business metrics that focus on tasks automated and hours saved, providing a glimpse into business impact. Finally, the third tier examines human factors, such as how employees and customers respond to the technology. The third tier is the most difficult for organisations to tackle at this point. A June 2025 systematic review published on arXiv, an online repository of scholarly papers that have not yet been published in academic journals, found human-centred measurements are used

Emerging frameworks suggest human factors – such as trust, adoption rates and willingness to delegate decisions to AI – should be measured alongside technical performance because systems often fail in deployment due to user resistance rather than technical flaws.

in only 30% of agentic AI projects. Emerging frameworks suggest human factors – such as trust, adoption rates and willingness to delegate decisions to AI – should be measured alongside technical performance because systems often fail in deployment due to user resistance rather than technical flaws.

“There’s the impact to the business vertical or mission vertical and the value that agentic solutions provide at that level,” says Sanghvi of Booz Allen Hamilton. “Then there is the workflow level. How did I actually improve this discrete workflow? And then there is the end-user level. How did I improve the individual user’s experience or job function?”

The evolution reflects broader changes in how IT demonstrates value. “Before DevOps, IT operations was a cost centre, just a budget item in the red,” says Jay Lyman, senior research analyst with the Cloud Native and Applied Infrastructure and DevOps channels at 451 Research. “There was a shift from total cost of ownership to return on investment. This infrastructure doesn’t have to be just a cost. It can be part of business benefits.”

Overcoming Cultural and Technical Barriers


Effective agentic IT operations depend as much on organisational readiness as on technology. Companies must address fundamental challenges in processes, culture

FIGURE 3

DevOps Teams Top Agentic Users

With almost 50% of use cases, they far outdistance cybersecurity teams

Percentage of IT operations use cases for agentic AI:



Source: Information Services Group, June 2025

and mindset if they hope to successfully implement agentic systems.

As Mohan says, “Technology is easy, processes are hard, people are impossible.” Many organisations struggle to accept autonomous systems making critical decisions. Mohan compares the situation to aviation: Fully automated take-offs exist, yet passengers resist pilotless flights. “Our mindset is like, ‘What do you mean a plane without a pilot? I am not flying in one,’” he says.

When employees resist the technology, internal company conflicts make the problem worse. “AI is actually a business initiative, not a technical initiative,” Mohan explains. “If you ask an IT person how quickly you can get AI enabled, they’ll say, ‘Wait, I can’t even tell you where to find the most authoritative source of data, and you want me to give you an AI-enabled application?’ IT departments are grudgingly being driven to become ready because businesses want AI, like, yesterday.”

This tension creates unintended consequences. “Departments are saying, ‘Look, if you’re not ready to support us, we will go rogue,’” Mohan says. “That’s the rise of shadow AI. People will copy data from their corporate laptop to their personal laptop and use ChatGPT while IT takes time to approve which models can or cannot be used.”

According to an April 2025 survey by London-based Ernst & Young Global Ltd., half (49%) of the 504 senior business leaders surveyed cite data privacy and security breaches as their top concern about agentic AI. This finding represents a 19-point increase since 2024.

In addition, cloud costs often exceed projections, leading companies to focus on cost reduction rather than new technology investments. Industry adoption patterns reflect these financial realities. “The banks, the finance groups, as always, are further along,” Linthicum says. “They have more money to spend on playing around with technology.” This leadership reflects both resources and necessity: Banks face intense pressure to reduce operational costs while maintaining strict security and compliance standards.

By contrast, sectors with tighter budgets are more reluctant to use the technology. “Manufacturing, retail – they’re not really even experimenting yet,” Linthicum says. “People who normally don’t spend a lot of money on technology are taking a wait-and-see approach on agentic AI.”

Organisational maturity plays a crucial role in determining the extent to which the technology is used. “The more mature the DevOps deployment and the deeper the organisation has integrated platform engineering – which means creating shared infrastructure and tools for development teams – the more likely they are to be early adopters of both generative AI and agentic systems,” 451 Research’s Lyman says.

The ISG survey confirms this situation. DevOps teams are leading IT operations’ agentic implementations in 48% of use cases, well ahead of cybersecurity (9%) and infrastructure management (8%). FIGURE 3

“ People don’t understand how they’ll fit into a world with agentic AI operations. We can focus on being innovative and creative instead of having our phones go off in the middle of the night to go into the data centre and kick a server. That’s no way to spend a career. ”

David Linthicum, independent consultant

While cybersecurity lags in deploying agents, the security teams that have adopted this technology show impressive results. A July 2025 global survey of 436 cybersecurity professionals by the International Information System Security Certification Consortium in Alexandria, Va., found that 30% of security teams have agents that automatically sort low- and medium-priority security alerts; these agentic AI systems group related warnings together and generate reports that explain to analysts why each alert matters and what should be done about it.

“Security and compliance are the low-hanging fruit for implementing gen AI and agentic AI into IT operations,” says Lyman. “AI can automatically determine whether code changes meet regulatory requirements, eliminating the guesswork for developers. The technology can process massive volumes of security data, compliance rules and system configurations simultaneously, which is something that would overwhelm human teams.”

Workforce Transition Strategies

The effective use of agentic IT operations requires deliberate workforce transition strategies that address both technical skills and cultural adaptation.

The most critical element involves reframing how employees perceive AI’s role in their work so they see it as an assistant rather than a replacement. This framework helps employees view agentic systems as tools for enhancement rather than threats to their job security.

Clear communication about organisational intentions can be essential for managing workforce concerns. “You have to be as clear as you can about the impact on people, because that is the biggest impediment to success,” says Thompson. Organisations should frame their approach explicitly, he adds, giving this example: “Our intention is to increase productivity by X% without reducing any jobs. We aim to free people from mundane, repetitive work so they can focus on more-interesting roles.”

Data supports this work augmentation approach rather than employees’ job replacement fears. Contrary to assumptions that AI reduces head count, 19% of organisations report that generative AI significantly increased staffing and 33% report a slight increase, according to the 451 Research survey. “This is counterintuitive to the idea that gen AI will write all of our software and deploy it,” Lyman explains. “We are not there yet.”

Skills gaps present ongoing challenges that require systematic attention. Configuration errors often stem from insufficient training, Linthicum explains. “It comes down to misconfiguration because [employees] were never trained to do it,” he says. “The suppliers didn’t train them, and the company didn’t either.” Enterprises must invest in comprehensive training programmes that prepare staff for new ways of working alongside agentic systems.

Linthicum adds a more optimistic perspective on workforce impact: “People should be excited about agentic AI,” he asserts. “Why be frightened? If anything, people don’t understand how they’ll fit into a world with agentic AI operations. We can focus on being innovative and creative instead of having our phones go off in the middle of the night to go into the data centre and kick a server. That’s no way to spend a career.”

The Future of Intelligent Operations

Agentic IT operations are expected to evolve beyond reactive and predictive functions toward automatically improving system performance across multiple cloud and on-premises platforms. The trajectory suggests increasingly sophisticated agent ecosystems capable of reasoning about complex interdependencies and making strategic decisions.

Cost optimisation presents a significant opportunity, particularly in hybrid and multicloud environments where pricing and availability constantly fluctuate. Forward-looking firms are experimenting with dynamic resource allocation, where systems automatically find available

“An agent should test multiple AI models and automatically select the most cost-effective option meeting accuracy requirements. This capability allows continuous performance and cost optimisation without human intervention.”

Sanjeev Mohan, owner, SanjMo

applications and inexpensive compute resources, running workloads as long as cost-effective capacity remains available. “An agent should test multiple AI models and automatically select the most cost-effective option meeting accuracy requirements,” Mohan says. “This capability allows continuous performance and cost optimisation without human intervention.” This approach requires that appropriate guardrails and controls are in place.

Industry analysts predict faster adoption of agentic AI. Gartner projects that by 2026, 30% of enterprises will have automated more than half of their network activities, up from less than 10% in mid-2023. Looking further ahead, the Stamford, Conn.-based consultancy forecasts that by 2028, 15% of day-to-day work decisions could be made autonomously by agentic AI, including IT tasks like incident triage, ticket handling and automated remediation.

The transformation may fundamentally reshape software development cycles. Traditional IT projects follow distinct phases: requirements gathering, design, development, testing, deployment and operations. Sanghvi predicts this linear approach will shift dramatically. “When do you go from plan to design to develop?” he asks. “Are they even discrete phases? Could all three happen in a single day?”

With AI agents handling much of the groundwork – analysing prior project data, sketching designs, gathering prototypes, running tests and deploying changes – these phases could collapse into rapid iterations where business requirements, system design and working code emerge simultaneously.

Conclusion

As IT departments move beyond pilot projects toward production deployments, the path forward requires balancing ambitious goals with practical realities.

The technology foundation is largely in place, but successful implementation demands more than technical capabilities. Organisations must invest simultaneously in multiple areas: comprehensive training programmes, robust security infrastructure, properly configured data pipelines and governance frameworks that treat agents like employees, with clear roles and boundaries.

For organisations willing to make that full commitment, the potential payoff extends beyond operational efficiency. Agentic systems possess a capacity for growth and learning that current automation lacks, enabling IT operations that continuously improve without constant human intervention through maintaining appropriate human oversight. As Linthicum puts it, “If you do this right, these agents should become the smartest things in the loop.”



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.