DW DIGITALISATION WORLD

ISSUE 2 2026

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

DIGITALISATIONWORLD.COM

FROM AUTOMATED TO ORCHESTRATED

EXAGRID®

Tiered Backup Storage

The future is here.
Tiered Backup Storage

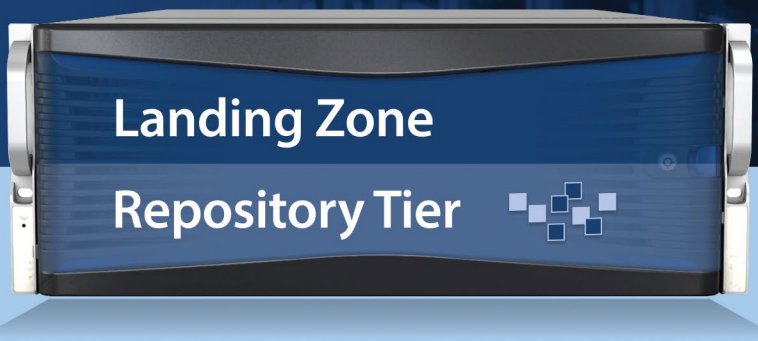

FASTEST BACKUPS

FASTEST RESTORES

SCALABILITY FOR FIXED-LENGTH BACKUP WINDOW

COMPREHENSIVE SECURITY WITH RANSOMWARE RECOVERY

LOW COST UP FRONT AND OVER TIME

MSP CHANNEL AWARDS 2025 WINNER

- BACKUP & DR INNOVATION OF THE YEAR
- STORAGE HARDWARE INNOVATION OF THE YEAR
- STORAGE VENDOR OF THE YEAR

Thank you so much to all who voted, and congratulations to our fellow MSP Channel Awards 2025 winners!

Visit our website to learn more about ExaGrid's award-winning Tiered Backup Storage.

LEARN MORE >

VIEWPOINT

BY PHIL ALSOP
EDITOR

AI amplifies the need for observability

OBSERVABILITY has long been treated as a technical discipline - something for engineers to worry about once systems are already in motion. But across this issue of Digitalisation World, a clearer narrative emerges: observability is no longer a supporting function. It is the foundation upon which AI-era operations either succeed or fail.

From Datadog's exploration of AI-powered SRE agents to Dynatrace's focus on log intelligence, the message is consistent: the challenge is no longer collecting data, but understanding it. Modern systems generate vast volumes of telemetry - logs, metrics, traces—but without context and correlation, more data simply creates more noise. The result is a paradox where organisations are data-rich yet insight-poor.

What's changing is the expectation of what observability should deliver. As highlighted by Nutanix, visibility that stops at infrastructure metrics is no longer enough. In AI-driven environments, observability must extend into model behaviour, data pipelines, and cost dynamics. It must answer not just "what happened," but "why it happened" and "what to do next." Without that, even the most advanced AI initiatives risk underperformance - or outright failure.

This is where AI itself enters the frame, not as a replacement for observability, but as its natural evolution. The concept of agentic AI, discussed by LogicMonitor, depends fundamentally on shared context and end-to-end visibility. Intelligent agents cannot coordinate effectively if they are operating on fragmented or incomplete data. In that sense, observability becomes the common language that allows systems, and increasingly, autonomous agents, to work together coherently.

Yet there is a subtle but important shift underway. Observability is moving from passive monitoring to active participation. AI SRE agents that investigate incidents, log platforms that prioritise meaningful signals, and systems that correlate behaviour across layers all point toward a more dynamic model. Observability is becoming an engine for decision-making, not just a source of information.

The implications for business are significant. Faster root-cause analysis, reduced downtime, and improved resilience are only part of the story. The real value lies in enabling organisations to operate with confidence at scale, where complexity is no longer a barrier but a managed variable.

If there is a unifying lesson from these perspectives, it is this: AI does not reduce the need for observability - it amplifies it. The more autonomous and distributed systems become, the more critical it is to see, understand, and act on what is happening beneath the surface.

In the end, observability is not about dashboards or data lakes. It is about clarity. And in an AI-driven world, clarity is competitive advantage.

Contents

Cover Story

From 'automated' to orchestrated: Closing the real productivity gap in the enterprise

Over the past two decades, enterprises have poured billions into automating business operations. Payroll runs itself. IT accounts are provisioned in seconds.

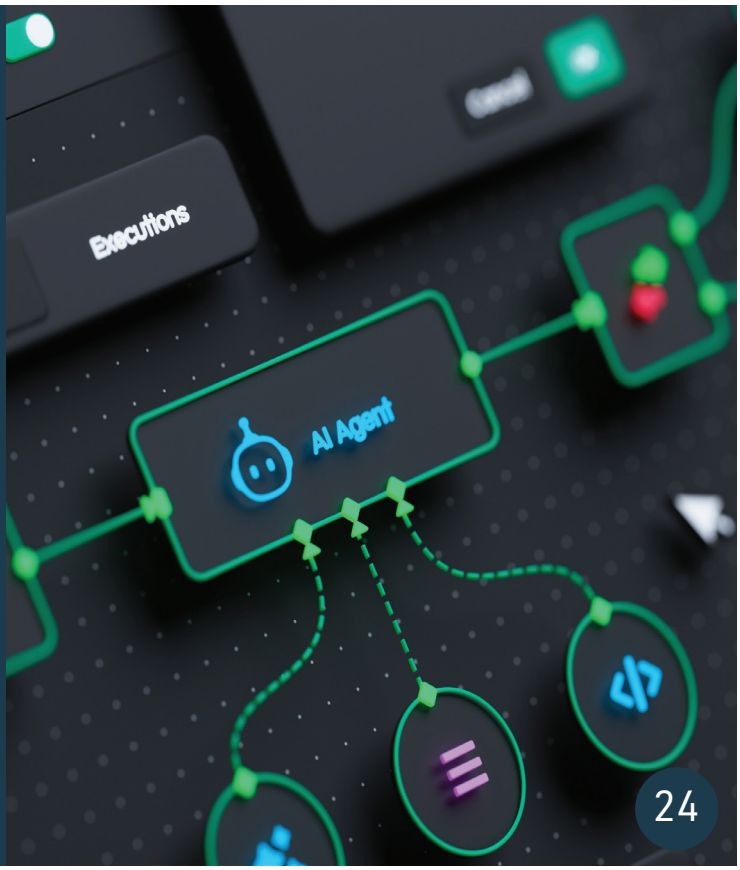

24

26 How to build an AI SRE agent that solves production incidents like a team of engineers

Site Reliability Engineers (SREs) are constantly flooded with alerts in large-scale, distributed environments, where every service, API, and infrastructure layer can fail.

28 Unlocking the true business value of modern log management

Recent outage highlights just how essential it is for businesses to have clear, real-time visibility into what's happening inside their digital systems.

26

30 Why agentic AI's next challenge is making systems work together

Agentic AI will continue to advance, that's a given. But its trajectory and benefit inside the enterprise will depend on how well organisations adapt.

32 What's observability? And why should I care if I've got AI?

Technology was supposed to make everything easier. Faster decisions, smarter systems, leaner operations. But for many leaders, the reality looks very different: rising costs, swelling cyber risk, and a tangle of legacy and multi-cloud complexity that's only getting harder to manage. AI is now promising to help solve this — but in truth, AI has its work cut out.

34 IT readiness in 2026 will be defined by data resilience - not digital ambition

For too long, organisations have mistaken digital acceleration for digital readiness. The rush to adopt AI, expand cloud usage and deploy new platforms has created an illusion of progress, but without resilient data foundations, much of that innovation remains fragile.

36 Ignore at your own risk: The hidden price tag of bad code

Bad code carries a cost that touches budgets, security, operations, employee wellbeing, and long-term competitiveness.

38 Why LLMs are plateauing – and what that means for software security

There's no doubt the AI-generated code landscape evolved at an unprecedented rate over the last year.

40 Why software is the essential building block behind quantum computing's huge potential

Quantum computing is no longer a distant concept.

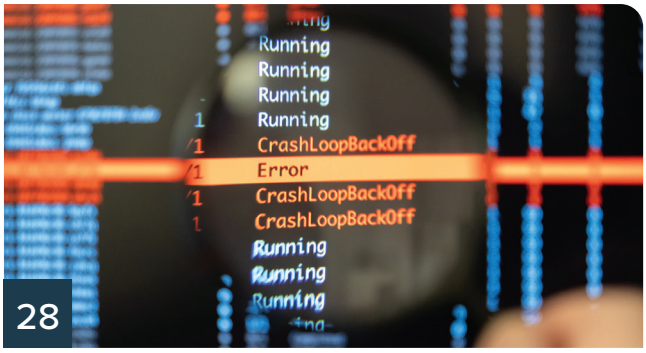

28

42 Unlocking AI's true potential: why high-quality first-party data and orchestration drive real business value

By investing in DEX platforms and solid data management tools, organisations can dodge the usual AI failures and actually tap into what it can do.

44 Does power consumption really offset the cost gap between Flash and HDD? Absolutely not

There is a persistent claim in the storage world that flash is worth its higher price as it consumes less power.

NEWS

06 The evolving role of CISOs in the AI era

07 Skill shortages: the roadblock to innovation in IT

08 Business concerns on post-pandemic AI prophecies

09 Unveiling the real impact of generative AI in the workplace

10 Data validation and governance: key focus for 2026

11 iManage report highlights gap between AI adoption and knowledge maturity

12 Navigating the challenges of AI: global report highlights systemic risks and governance gaps

13 The impact of Agentic AI on the future of customer service


12

DIGITALISATION WORLD

Editor
Philip Alsop
+44 (0)7786 084559
philip.alsop@angelbc.com

Senior B2B Event & Media Executive
Mark Hinds
+44 (0)2476 718971
mark.hinds@angelbc.com

Director of Logistics
Sharon Cowley
+44 (0)1923 690200
sharon.cowley@angelbc.com

Design & Production Manager
Mitch Gaynor
+44 (0)1923 690214
mitch.gaynor@angelbc.com

Graphic Design & Multimedia Assistant
Harvey Watkins
harvey.watkins@angelbc.com

Publisher
Jackie Cannon
+44 (0)1923 690215
jackie.cannon@angelbc.com

Circulation & Subscriptions
+44 (0)1923 690214
circ@angelbc.com

Directors
Scott Adams: CTO
Sukhi Bhadal: CEO

Published by: Angel Business Communications Ltd, 6 Bow Court, Burnsall Road, Coventry CV5 6SP
T: +44 (0)2476 718970 E: info@angelbc.com

Angel BUSINESS COMMUNICATIONS

Digitalisation World is published 10 times a year on a controlled circulation basis in Europe, Middle East and Africa only. Subscription rates on request. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd.

ISSN 2396-9016 (Online)

The evolving role of CISOs in the AI era

A survey of 650 global CISOs examines how security leaders are navigating AI adoption, expanding responsibilities, workforce challenges, and cross-organisational collaboration in today's evolving threat landscape.

CISCO has announced the release of Splunk's annual report, The CISO Report: From Risk to Resilience in the AI Era, based on a survey of 650 global Chief Information Security Officers (CISOs). The report explores the expanding responsibilities of CISOs and their approach to AI adoption, workforce investment, and risk management in a complex security landscape.

The findings indicate that AI is increasingly viewed as an important capability for security teams. Key insights from the survey include:

- 95% of CISOs identify the growing sophistication of threat actors as their primary risk.
- 92% prioritise improvements in threat detection and response, identity and access management, and investment in AI-based cybersecurity capabilities.
- 89% report that AI enhances data correlation, supporting improved incident visibility.
- 82% say AI contributes to faster data analysis and response.

At the same time, adoption of AI brings concerns: 86% of CISOs believe AI could increase the sophistication of social engineering attacks, and 82% are concerned it may add complexity and accelerate deployment challenges.

The report notes that CISOs are operating with expanded responsibilities during digital transformation, with more than three-quarters expressing concern about personal accountability for security incidents. Responsibilities increasingly include AI governance and oversight, alongside secure software development (DevSecOps).

Despite automation advances, human expertise remains central to security strategy. Organisations report prioritising workforce upskilling, hiring, and contractor support to address skills gaps and maintain oversight.

Collaboration and shared accountability across departments are highlighted as important factors for improving cybersecurity outcomes. Many respondents emphasise the value of aligning security initiatives with broader organisational objectives.

The report notes that CISOs are operating with expanded responsibilities during digital transformation, with more than three-quarters expressing concern about personal accountability for security incidents. Responsibilities increasingly include AI governance and oversight, alongside secure software development (DevSecOps).

Workforce challenges remain significant, with high alert volumes and tool fatigue contributing to stress. Strategies such as consolidating security data and communicating security metrics in business terms are used to support internal alignment and reduce operational pressure.

Overall, the report outlines how CISOs are adapting to evolving risks by integrating AI thoughtfully, strengthening governance, and positioning cybersecurity as a contributor to broader business resilience.

Skill shortages: the roadblock to innovation in IT

Skill shortages are stifling innovation in IT, with firms focusing on talent acquisition and retention to drive future growth.

ACCORDING TO RECENT research by Ayming, an innovation consultancy, the largest barrier to innovation within IT industries is a shortage of skilled talent. In a survey of IT and technology business leaders, 41% identified this skills gap as the primary hurdle. A further 35% cited unstable research and development (R&D) budgets as a significant challenge.

In response, a quarter of IT firms are prioritising talent attraction and retention as part of their strategic planning towards the end of 2026. Addressing these skills shortages is viewed as essential to sustaining innovation activity.

The survey also identified a growing trend towards offshore R&D. Among the 56% of IT businesses that outsource R&D, 40% stated that improved access to the necessary talent was a key reason. In addition, 26% indicated that lower wage costs for R&D talent were an important factor.

Mark Smith, Managing Director at Ayming UK, commented: "When we talk about innovation, the focus often falls on shiny new technologies, but the reality is that innovation is driven by people. It's people who develop new technologies, apply them in the real world, and carry out the R&D that turns ideas into commercial success."

Smith added: "The IT sector is already innovating as pace, but sustaining that momentum depends on access to the right specialist skills. Without a sustainable pipeline of talent, there is real risk that skills shortages will push more R&D and development activity overseas, leaving firms with the outcomes of innovation, rather than the infrastructure needed to build and deliver it in the long term."

Key findings include:

- Challenges: Talent shortages are identified as the top barrier to technological innovation, with unstable R&D budgets following closely behind.
- Investment: Despite ongoing economic pressures, 90% of technology firms intend to increase or significantly increase their innovation budgets in 2026.
- Funding sources: Self-funding remains the most common approach to financing innovation projects, followed by equity funding (46%) and innovation grants (44%).
- Strategic outlook: More than half of IT and technology companies now have innovation strategies that extend over the next decade, exceeding levels seen in other sectors.

The growing divide: security struggles to keep up with software development

IN ITS 2026 State of Software Security Report, Veracode, a global application risk management provider, highlights a widening gap between software development speed and security efforts. The report shows that 82% of organisations are dealing with security debt — an increase of 11% compared with the previous year.

Of those organisations, 60% are categorised as having "critical" security debt, meaning accumulated vulnerabilities that could cause significant damage if exploited. To address this, the report recommends adopting a "Protect, Prioritize, and Prove" approach to reduce risk in 2026 and beyond.

Now in its 16th edition, the report analysed more than 1.6 million unique applications across enterprises, commercial software suppliers, software outsourcing providers, and open-source projects globally. It identifies a clear imbalance between rapid development cycles and the pace at which vulnerabilities are remediated.

While detection capabilities have improved, unresolved vulnerabilities continue to accumulate. High-risk vulnerabilities have increased by 36% year-over-year, defined as flaws that are both severe and highly exploitable.

The findings suggest that high-risk vulnerabilities require stronger prioritisation, moving beyond generic severity scoring toward assessments based on real-world attack potential. Security debt is also influenced by greater reliance on open-source components, which account for 66% of the most persistent vulnerabilities.

To reduce these risks, Veracode recommends a strategic framework centred on Prioritize, Protect, and Prove, enabling organisations to focus on safeguarding their most critical systems and applications that hold essential data.

The report also notes the impact of AI on the landscape, introducing new high-risk vulnerability patterns while AI-driven remediation tools begin to offer additional support in closing gaps.

As organisations manage growing security debt, the emphasis is on prioritising the most significant risks rather than attempting to eliminate every vulnerability, while maintaining alignment with security and compliance requirements.

Business concerns on post-pandemic AI prophecies

Large enterprises express concern that AI may not deliver the resilience and business continuity benefits expected, with cybersecurity identified as the leading threat to operations.

ACCORDING TO recent findings from supply chain intelligence firm Zero100, large enterprises are increasingly cautious about disruption while expressing limited confidence in AI delivering the resilience it promises. Despite public commitments, views differ significantly when assessing AI's contribution to operational strength.

In a survey of Chief Operating Officers (COOs) from companies valued at over $1bn, cybersecurity emerged as the primary concern for business continuity. Expectations around AI-driven transformation remain measured and are not fully aligned with the messaging presented by senior executives to shareholders.

A total of 35% of businesses identify cyber incidents as the main risk to continuity over the next year, ahead of geopolitical instability (20%), trade policy shocks (16%) and labour disruption (8%). Cyber incidents are not only the leading concern but are also perceived as the fastest-moving threat. Nearly two-thirds (62%) believe they can respond to a cyber incident within minutes or hours, compared with disruptions such as tariffs, which require days or weeks for the majority (83%) of businesses to respond.

There is mixed opinion among COOs regarding AI's impact on cyber risk — around half (50%) believe it will improve resilience, while 43% believe it could increase vulnerabilities. However, most expect AI to support supply chain management (64%) and help address skills shortages (58%).

The research also highlights a gap between external communications and internal confidence. Fewer than one in five COOs (17%) believe that the majority of their company's AI commitments to investors can be delivered on time.

Although AI is positioned by CEOs as a productivity driver, internal sentiment is more cautious. Operations leaders recognise the potential of AI but question the realism of projected timelines. Expectations around agentic AI remain restrained, with only 7% believing such systems will fundamentally redesign most workflows within two years. The most common view, held by 43% of COOs, is that agentic AI will affect between 11% and 25% of workflows, indicating targeted adoption rather than broad transformation.

Progress in deploying agentic AI at scale appears to be constrained more by organisational readiness than by technical capability. When asked to assess preparedness for large-scale implementation, COOs rated technology infrastructure highest at 6.2 out of 10, followed by leadership understanding (6.0), data foundations (5.8), process maturity (5.6) and workforce skills (5.5).

Despite the potential for AI-enabled transformation, most organisations continue to prioritise cost control in their day-to-day operations. Cost management accounts for nearly a third of COO performance metrics on average (29%), compared with revenue growth (15%) and customer service improvement (17%). Longer-term considerations such as brand trust represent a smaller proportion, typically around 8%.

Lauren Acoba, VP, Research & Advisory Services at Zero100, commented that organisations have access to the necessary tools but lack the internal alignment required to reshape workflows through agentic AI. She stated: "Companies have the tools, but not yet the organisational muscle to change how work gets done using agentic AI," Acoba added. "Cost control still dominates how they're measured; they're asked to chase growth, but are judged on profit. The result is an operating model that rewards not breaking things, more than trying new ideas. Until that changes, AI's impact on resilience will be more incremental than transformational."

Unveiling the real impact of generative AI in the workplace

Despite significant investment in GenAI, employee engagement remains limited, spotlighting challenges in integration and adoption.

IN RECENT YEARS, enterprises have invested as much as $40 billion into Generative AI (GenAI). However, research from Nexthink indicates significant underutilisation, with most employee interactions with GenAI tools lasting under four minutes on average.

The analysis, which covers 4.9 million sessions daily and incorporates input from 3.4 million employees, provides insight into how frequently these technologies are used in practice. Employees average around ten interactions per day, yet total weekly engagement amounts to approximately three hours and fourteen minutes — around thirty-nine minutes per day.

The findings suggest a pattern of short "micro sessions" rather than sustained integration into everyday workflows.

Despite limited engagement, the analysis shows users save approximately three hours and forty-seven minutes per week on average through the use of GenAI tools. However, performance varies across the four leading tools in the market:

- ChatGPT: Average engagement — 2 hours 47 minutes; Net time saved — 5 hours 46 minutes
- Claude: Average engagement — 2 hours 30 minutes; Net time saved — 3 hours 23 minutes
- Copilot: Average engagement — 2 hours 40 minutes; Net time saved — 2 hours 45 minutes
- Gemini: Average engagement — 2 hours 13 minutes; Net time saved — 4 hours 46 minutes

The differences indicate uneven adoption and efficiency outcomes across platforms, suggesting scope for more consistent deployment and optimisation.

While organisations have adopted these tools at scale, limited visibility into usage — including who is using them and for what purposes — can restrict understanding of how value is being derived from the investment. Nexthink's AI Drive platform aims to address this by consolidating data on usage, measurement and guidance to provide a more integrated view of AI activity within organisations.

To maximise returns on GenAI investment, organisations may need to move beyond providing access alone and focus on structured adoption strategies. Understanding how each tool contributes across teams and use cases can help inform targeted support and training.

Encouraging a workplace culture that supports experimentation and views AI as a mechanism for improving processes and adapting existing approaches may also support deeper integration. Ongoing training and regular evaluation of tool effectiveness can further strengthen adoption and alignment with operational needs.

For organisations seeking to increase the impact of GenAI, developing clearer visibility of employee requirements and operational challenges remains important. Insights from usage analytics platforms can assist in identifying adoption gaps and improving integration into everyday workflows.

The effectiveness of GenAI in the workplace will depend largely on how it is implemented and embedded. While investment has established the foundation, outcomes will be shaped by adoption practices and governance over time.

To maximise returns on GenAI investment, organisations may need to move beyond providing access alone and focus on structured adoption strategies. Understanding how each tool contributes across teams and use cases can help inform targeted support and training.

Data validation and governance: key focus for 2026

The Workiva 2026 Benchmark Survey highlights data validation and governance as top priorities for businesses, crucial for financial and non-financial reporting.

FOR YEARS, corporate leaders have been advised about the consequences of poor data quality. As organisations expand their use of artificial intelligence, addressing this issue is becoming increasingly important. Workiva's 2026 Benchmark Survey Report, which surveyed nearly 1,500 executives across various departments, highlights these challenges.

The survey indicates that 62 percent of UK respondents have prioritised automating data collection and validation within their digital transformation agendas for 2026. A stronger focus on data governance is also reflected, with 20 percent identifying it as a key area.

Many organisations report allocating dedicated budgets and IT support to these initiatives. According to the survey, 84 percent are assigning specific funding, while 80 percent have involvement from IT teams in supporting these efforts.

Data fragmentation is identified as an ongoing concern. Without a consolidated view of data, business leaders may lack access to insights that inform decision-making, which can create operational and strategic risks. Inadequate data management may also contribute to regulatory and compliance challenges.

The survey shows that 44 percent of UK-based respondents experience barriers at work related to data issues. These include a lack of real-time data (20 percent) and restricted access across departments (24 percent). Such challenges can limit the role of data in decision-making processes.

Some organisations are investing in integrating disparate data systems and developing more unified reporting approaches. Efforts often include improving data quality and governance and enhancing collaboration across finance, technology and sustainability teams.

By consolidating data sources and standardising reporting, businesses aim to reduce risks associated with fragmented information and better manage potential challenges. This approach supports regulatory compliance, operational consistency and data-informed decision-making.

With data, AI and automation increasingly interconnected, organisations are considering unified data strategies as part of broader business planning. Establishing consistent data standards can help support operational visibility and informed growth as business environments evolve.

The rising tide of ransomware: unveiling 2025's alarming trends

RANSOMWARE remains a persistent threat, with the latest 2025 State of Ransomware Report by BlackFog unveiling a 49% increase in ransomware incidents globally. This highlights the evolving nature of cyber threats, now powered notably by AI-driven techniques.

In 2025, ransomware tactics evolved, favouring stealth and speed over traditional disruptive approaches. As a result, the year witnessed 1,174 publicly disclosed cyber incidents, a peak since 2020. Furthermore, a 37% increase in attacks went undisclosed, underscoring the clandestine nature of modern ransomware strategies.

A notable revelation from the report is the deployment of large-scale, AI-enabled cyber attacks. This shift reflects an evolution in ransomware activity, exemplified by incidents such as the reported misuse of AI models to conduct autonomous reconnaissance, exploitation, and data theft.

Such threats underscore the challenge defenders face, requiring countermeasures for effective data protection.

Among the various industries, retail and healthcare sectors remain prime targets. In 2025, the healthcare sector accounted for 22% of all ransomware attacks, maintaining its status as the most targeted industry. The retail sector, too, saw a spike, with high-profile brands experiencing attacks, as cybercriminals aim for rich databases housing valuable consumer data.

Across the globe, the prevalence of attacks echoes a grim reality: no nation is immune. Organisations in 135 countries—69% globally—faced ransomware threats in 2025. Predominantly, the United States was the most affected region, accounting for a significant 58% of all recorded attacks.

The surge in undisclosed incidents highlights the necessity for more transparent reporting and enhanced security measures. As AI-enabled threats become common, organisations must prioritise strategies to prevent data exfiltration, the core goal of modern ransomware groups.

BlackFog's ADX technology is positioned as a tool designed to help organisations detect and mitigate data breaches and ransomware risks. As threats continue to adapt, a concerted effort towards advanced cybersecurity solutions is essential to protect vital infrastructure and sensitive data.

iManage report highlights gap between AI adoption and knowledge maturity

iManage's latest report examines the relationship between knowledge management maturity and AI adoption.

IN TODAY'S fast-paced digital environment, organisations are working to integrate artificial intelligence (AI) into their daily operations. However, a recent iManage report indicates that only a minority have effectively embedded AI into their workflows, despite widespread interest in AI initiatives. The research, based on responses from 3,185 professionals across 26 countries, highlights the gap between AI ambition and practical implementation and underscores the role of knowledge foundations.

Published by iManage, the global study examines the current state of AI maturity among professional services firms. The findings show that organisations with mature and well-governed knowledge foundations demonstrate higher levels of AI adoption and operational performance compared to others.

According to the report, 85% of firms are either piloting, implementing, or actively using AI. However, only 17% have integrated AI into their regular operations. This gap suggests that access to AI tools alone does not ensure consistent deployment. Effective use of AI depends on the strength of an organisation's knowledge management capabilities.

85% of firms are either piloting, implementing, or actively using AI. However, only 17% have integrated AI into their regular operations.

The report indicates that organisations with higher maturity in knowledge work report stronger business outcomes. Companies with more developed knowledge systems are nearly twice as likely to report year-on-year growth, improved profitability, and better financial performance.

These knowledge-mature organisations are also deploying AI within both operational and client-facing workflows. When customer demand informs AI strategies, firms may be better positioned to respond to changing expectations.

One of the challenges identified in the report is the impact of governance gaps. Many firms experience policy-related incidents due to the use of unregulated AI tools, and some delay broader AI adoption because of security concerns. The findings point to the importance of defined governance frameworks that aim to protect operations while supporting AI integration.

Ongoing investment in document and knowledge management systems reflects an effort to address these governance concerns. While 72% of organisations plan further investment within the next two years, the report notes that outcomes depend on the quality and governance of systems rather than technology alone.

The report highlights the importance of knowledge architecture and structured AI adoption for organisations seeking to remain competitive. As AI moves from experimentation toward broader integration, organisations that develop structured knowledge practices may be better positioned to support implementation.

The iManage Knowledge Work 2026 Benchmark Report provides insight into current practices and trends. It reflects ongoing efforts to combine knowledge management with AI capabilities to support operational and business objectives.

In conclusion, organisations that maintain a structured and well-governed knowledge framework may be better placed to translate AI experimentation into longer-term operational impact.

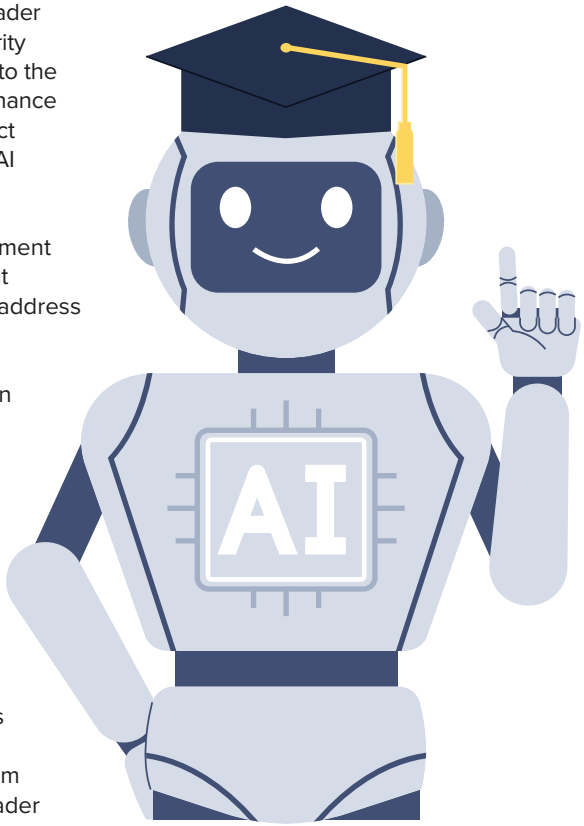

Navigating the challenges of AI: global report highlights systemic risks and governance gaps

The International AI Safety Report advocates for strengthened AI governance and highlights potential risks related to misuse and cognitive offloading.

THE SECOND INTERNATIONAL AI Safety Report has been released, examining the risks and challenges associated with the rapid development of artificial intelligence systems. Experts from institutions worldwide collaborated to evaluate evolving threats and assess the implications of advanced AI technologies.

The report notes that the rapid progression of general-purpose AI models presents both opportunities and challenges. As capabilities in reasoning, autonomy, and multimodal functions increase, concerns arise around misuse, systemic risks, misinformation, cybersecurity vulnerabilities, and reduced human oversight.

A particular focus is placed on systemic risks across critical domains. Improperly managed AI systems can create regulatory, reputational, and operational vulnerabilities, especially within financial services. Differences in governance standards internationally may further increase exposure, potentially allowing exploitation by malicious actors.

Financial systems that use AI for onboarding, transaction monitoring, or fraud detection need transparency and accountability in deployment. Aligning safety principles with practical implementation is essential, including clear standards for explainability, auditability, and human oversight to ensure responsible AI use.

The report also highlights the trend of 'cognitive offloading', where human decision-making is increasingly delegated to AI systems. While this can improve efficiency, there is a potential risk to critical thinking skills and institutional expertise over time.

To mitigate these issues, the report recommends enhanced international collaboration, increased transparency from AI developers, and thorough safety testing. Key areas examined include capability development, misuse risks, systemic impacts, and governance gaps, emphasising the need to align global AI governance with technological advances.

Analysts also stress the importance of strong data quality and governance frameworks, particularly in financial services. Reliable, high-volume, multi-source data pipelines are critical for supporting AI-driven decision-making. Strengthened governance and international cooperation can help balance competitiveness with caution, enabling the benefits of AI while addressing emerging risks.

In summary, the report underscores that as AI continues to expand across sectors, cross-border governance, safety standards, and data management practices are essential for mitigating risks and supporting responsible adoption.

The impact of Agentic AI on the future of customer service

IN TODAY'S business landscape, the integration of Agentic AI represents a shift in customer service operations. This form of artificial intelligence goes beyond traditional automation by adapting to customer needs and influencing how interactions are managed. As enterprises adopt these technologies at scale, understanding their capabilities and implications remains important.

Agentic AI refers to AI systems designed to achieve predefined outcomes by understanding, planning, and executing required actions autonomously. Unlike earlier technologies that followed predefined scripts, these systems adapt to real-time inputs, enabling more personalised service delivery.

The Agentic AI CX Frontline report highlights early adopters reporting faster deployment cycles, in some cases lasting only a few weeks, and containment rates of up to 80%. Metrics such as Customer Satisfaction Score (CSAT) have reportedly improved by up to 20%.

The adoption of this approach affects traditional customer service structures. By adjusting operational processes, organisations may reduce the cost per contact and modify workforce models. This often involves shifting human agents away from routine tasks toward roles that require more complex decision-making and oversight, increasing the level of human involvement in the AI-driven workflow.

Implementing Agentic AI typically requires several considerations:

- Data Integration: A consistent flow of accurate data is necessary for AI systems to operate effectively.
- Organisational Readiness: Companies need technical and cultural preparedness for integration, including role adjustments and skills development.
- Governance and Trust: Clear governance frameworks are required to address concerns around reliability and operational use.

Enterprises adopting this approach aim to build integrated engagement environments where AI systems support interactions alongside existing operational structures. This approach is intended to improve customer experience while supporting operational efficiency.

As AI becomes more embedded in customer service operations, success depends on integration strategy, planning, and governance. Agentic AI is increasingly positioned as a component of evolving customer service models.

Harnessing AI alignment: unlocking billions in workplace value

ARTIFICIAL INTELLIGENCE (AI) has significant potential to enhance workplace productivity and efficiency. A recent report from Zellis examines how improved AI alignment between leaders and employees could generate economic and cultural benefits for UK businesses.

Zellis' research identifies a "grey zone" of AI misalignment, where the strategic use of AI intended by leaders does not always reflect employees' day-to-day experiences. While 94% of business leaders confirm AI tool usage within their organisations, 61% of employees report regular interaction with these tools. This difference highlights a gap between strategy and practical application, limiting the potential impact of AI in the workplace.

The report suggests that reducing this misalignment could lead to substantial economic effects. Employees and leaders indicate that better alignment between AI usage and work processes could recover approximately £40 billion in staff time value annually. This represents around 1.7 billion hours of working time that could potentially be redirected towards other tasks.

Improved AI alignment is also associated with potential reductions in operational costs of up to £20 billion. Leaders estimate that enhancing alignment could reduce ongoing operational costs by up to 10%, creating opportunities for reinvestment within organisations.

Differences in adoption across age groups are evident. Gen Z and Millennial employees are more likely to use AI tools regularly compared to Gen X and Baby Boomer employees. This generational variation points to a need for training and education to support broader engagement across different employee groups.

Integrating AI in a way that supports human decision-making is highlighted as important. Organisations that implement transparent AI strategies and maintain clear communication may strengthen trust and engagement. Around 61% of employees who use AI tools report that it improves their work quality and productivity.

To maximise the potential of AI, the focus is placed on alignment between strategy and implementation rather than adoption alone. AI functions as a tool to support human capability and decision-making, requiring ongoing communication and structured engagement with teams. Encouraging feedback and supporting skill development may help organisations make better use of AI capabilities.

AI alignment represents both potential economic impact and an opportunity to support workplace engagement through clearer integration of technology into business processes.

Gartner identifies the top cybersecurity trends for 2026

The chaotic rise of AI, geopolitical tensions, regulatory volatility and an accelerating threat landscape are the driving forces behind the top cybersecurity trends for 2026, according to Gartner.

"CYBERSECURITY LEADERS are navigating uncharted territory this year as these forces converge, testing the limits of their teams in an environment defined by constant change," said Alex Michaels, Director Analyst at Gartner. "This demands new approaches to cyber risk management, resilience and resource allocation."

The following six trends will have broad impact across transforming governance, securing new frontiers and normalizing AI adoption.

Trend 1: Agentic AI demands cybersecurity oversight

Agentic AI is rapidly being used by employees and developers, creating new attack surfaces. No-code/low-code platforms and vibe coding expand this further, driving unmanaged AI agent proliferation, unsecured code and potential regulatory compliance violations.

"While AI agents and automation tools are becoming increasingly accessible and practical for organizations to adopt, strong governance remains essential," said Michaels. "Cybersecurity leaders must identify both sanctioned and unsanctioned AI agents, enforce robust controls for each and develop incident response playbooks to address potential risks."

Trend 2: Global regulatory volatility drives cyber resilience efforts

Shifting geopolitical landscapes and evolving global mandates have made cybersecurity a critical business risk with direct implications for organizational resilience. With regulators increasingly holding boards and executives liable for compliance failures, inaction can result in substantial penalties, lost business and irreversible reputational damage.

Gartner advises cybersecurity leaders to formalize collaboration across legal, business and procurement teams to establish clear accountability for cyber risk. Aligning control frameworks to recognized standards and addressing data sovereignty concerns will help reduce compliance gaps.

Trend 3: Postquantum computing moves into action plans

Gartner predicts advances in quantum computing will render the asymmetric cryptography organizations rely on to secure data and systems unsafe by 2030. Postquantum cryptography alternatives must be adopted now to avoid potential data breaches, legal liability and financial loss from "harvest now, decrypt later" attacks targeting long-term sensitive data.

"Postquantum cryptography is reshaping cybersecurity strategies by prompting organizations to identify, manage and replace traditional encryption methods, while prioritizing cryptographic agility," said Michaels. "By investing in these capabilities and prioritizing migration now, assets will be secured when quantum threats become a reality."

Trend 4: Identity and access management adapts to AI agents

The rise of AI agents is introducing new challenges to traditional identity and access management (IAM) strategies, especially in identity registration and governance, credential automation and policy-driven authorization for machine actors. Failure to address these issues will lead to greater risk of access-related cybersecurity incidents as autonomous agents become more prevalent.

Gartner recommends taking a targeted, risk-based approach, by investing where gaps and risks are greatest while leveraging automation where possible. This is essential for enabling innovation, ensuring compliance and securing critical assets in AI-centric environments.

"To realize the full potential of AI in security operations, cybersecurity leaders must prioritize people as much as technology." – Alex Michaels, Director Analyst at Gartner

Trend 5: AI-Driven SOC solutions destabilize operational norms

Driven by cost optimization practices and increasing interest in AI, the emergence of AI-enabled security operations centers (SOCs) is introducing new complexity. This is contributing to staffing pressures, increased upskilling demands and evolving cost considerations for AI tools, even as these technologies enhance alert triage and investigation workflows.

"To realize the full potential of AI in security operations, cybersecurity leaders must prioritize people as much as technology," said Michaels. "Strengthening workforce capabilities, implementing human-in-the-loop frameworks into AI-supported processes and aligning adoption with clear strategic objectives will be critical to maintaining resilience as SOCs evolve."

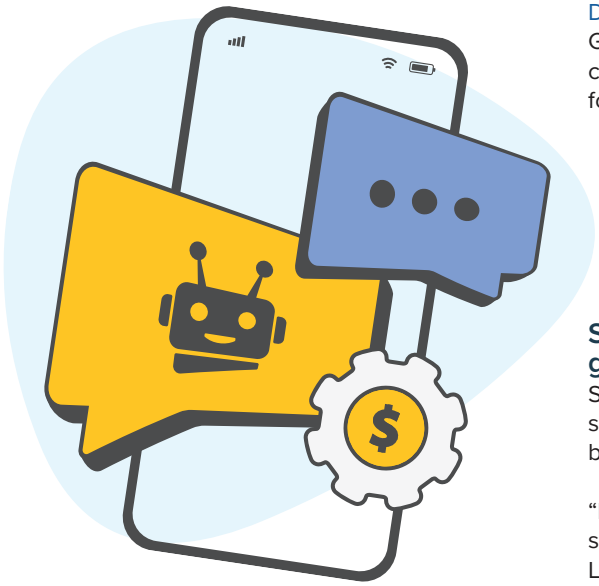

	2025 Spending	2025 Growth (%)	2026 Spending	2026 Growth (%)
Data Center Systems	496,231	48.9	653,403	31.7
Devices	788,335	9.1	836,417	6.1
Software	1,249,509	11.5	1,433,633	14.7
IT Services	1,717,590	6.4	1,866,856	8.7
Communications Services	1,303,651	3.8	1,365,184	4.7
Overall IT	**5,555,316**	**10.3**	**6,155,493**	**10.8**

➤ Table 1. Worldwide IT Spending Forecast (Millions of U.S. Dollars)

Trend 6: GenAI breaks traditional cybersecurity awareness tactics

Existing security awareness efforts continue to fail to reduce cybersecurity risks as GenAI adoption accelerates. A Gartner survey of 175 employees conducted between May and November 2025 indicates over 57% use personal GenAI accounts for work purposes and 33% admit inputting sensitive information into unapproved tools.

Gartner recommends shifting from general awareness training to adaptive behavioral and training programs that include AI-specific tasks. Strengthening governance, embedding secure practices and establishing policies for authorized use will reduce exposure to privacy breaches and intellectual property loss.

Worldwide IT spending to grow 10.8% in 2026

Worldwide IT spending is expected to reach $6.15 trillion in 2026, up 10.8% from 2025, according to the latest forecast by Gartner.

"AI infrastructure growth remains rapid despite concerns about an AI bubble, with spending rising across AIrelated hardware and software," said John-David Lovelock, Distinguished VP Analyst at Gartner. "Demand from hyperscale cloud providers continues to drive investment in servers optimized for AI workloads."

Server spending is projected to accelerate in 2026, growing 36.9% year-over-year. Total data center spending is expected to increase 31.7%, surpassing $650 billion in 2026, up from nearly $500 billion the previous year (see Table 1).

Software spending shows second-highest growth potential despite lower revision

Software spending growth for 2026 has been slightly revised downward to 14.7%, from 15.2% for both application and infrastructure software.

"Despite the modest revision, total software spending will remain above $1.4 trillion," said Lovelock. "Projections for generative AI (GenAI) model spending in 2026 remain unchanged, with

Geography	2025	2026	2027
China (Region)	37,539	47,379	58,544
North America	12,667	16,394	21,127
Europe	6,868	12,587	23,118
Mature Asia/Pacific	851	1,593	3,155
Japan Region	519	932	1,816
Emerging Asia/Pacific	430	755	1,326
Latin America	278	506	946
Middle East and North Africa	132	250	515
Sub-Saharan Africa	16	31	61
Total	**59,300**	**80,427**	**110,609**

Table 2. Sovereign Cloud IaaS Spending by Region, 2025-2027 (Millions of U.S. Dollars)

growth expected at 80.8%. GenAI models continue to experience strong growth, and their share of the software market is expected to rise by 1.8% in 2026."

Device growth expected to slow in 2026

Shipments of mobile phones, PCs, and tablets continue to grow steadily. Total spending on devices is projected to reach $836 billion in 2026. However, market-demand constraints will slow growth to 6.1% in 2026.

"This slowdown is largely due to rising memory prices, which are increasing average selling prices and discouraging device replacements," said Lovelock. "Additionally, higher memory costs are causing shortages in the lower end of the market, where profit margins are thinner. These factors are contributing to more muted growth in device shipments."

Three pillars for deriving value from AI

Only one out of every five data and analytic (D&A) or AI leaders are concerned that uncertain costs will limit AI value according to Gartner.

A Gartner survey of 353 D&A and AI leaders from November through December 2025 found that this has led to only 44% of organizations adopting financial guardrails or AI FinOPs practices.

"Where adoption rates for AI deployment have grown from just two out of five organizations in 2024, to four out of five organizations today, D&A leaders must achieve clarity and focus on ROI to better achieve the growing AI goals and ambitions of their organizations," said Adam Ronthal, VP Analyst at Gartner. "D&A leaders must realize they are responsible for delivering real value in the midst of all this AI hype and fears of an AI bubble that might burst."

"Getting to value is often measured using ROI, which D&A leaders need to think of as more than just a financial measure," said Georgia O'Callaghan, Director Analyst at Gartner. "There are three ways to approach value that will help D&A leaders steer their organizations safely and effectively through the turbulent AI value waters."

During the opening keynote at the Gartner Data & Analytics Summit, taking place here through Wednesday, Gartner analysts discussed these three ways to derive value from AI.

Set AI ambition

Increased acceleration and uncertainty, combined with concerns about trust and control, drive the need for continuous learning and adaptation.

"D&A leaders may be experimenting with AI and learning a lot, but that also means they risk falling behind because everyone is experimenting," said Ronthal. "D&A leaders should set their AI-ambition to help them maximize value from the insights their data provides, together with the knowledge and intuition of their team. This provides a return on intelligence."

To set this level of ambition, D&A leaders must radically rethink the impact of AI on D&A, set a shared vision and determine their level of AI ambition, take AI leadership, decide their role and manage the unpredictable and hidden costs of AI early.

Strengthen AI foundations

Without strong foundations, AI will remain what it is for most organizations today; an expensive experiment.

"Expecting AI or GenAI to compensate for delayed upgrades, siloed teams and years of technical debt is wishful thinking," said O'Callaghan. "D&A leaders must make sure their data is AI-ready, prevent exposing the wrong data to the wrong people and avoid inaccuracies, misunderstandings and hallucinations with a well-designed context layer. This provides a return on integrity."

To create strong AI foundations and reduce risk, D&A leaders should align their foundational initiatives with their AI ambition level, make governance a value accelerator and create a single, unified context layer.

Empower people for AI transformation

While organizations change at a rapid pace, humans have a finite capacity to incorporate change. AI readiness grows much faster than human readiness.

"D&A leaders must make the shift from thinking about roles to focusing on skills with respect to AI," said Ronthal. "D&A leaders will get value from their investments in developing their workforce. By focusing on skills, mindset, and behavioral change, they can unlock both individual and collective potential. This will increase employee engagement and productivity, making their organization more adaptive to change. Ultimately, this provides a return on individuals."

To empower people for AI-driven transformation, D&A leaders must substantially budget for change

management, prioritize mindset and skillset over toolset, address employee concerns with a skills-development roadmap and also pilot fusion teams of blended human and artificial intelligence.

Worldwide sovereign cloud IaaS spending to total $80 billion

Worldwide sovereign cloud infrastructure as a service (IaaS) spending is forecast to total $80 billion in 2026, a 35.6% increase from 2025, according to Gartner.

"As geopolitical tensions rise, organizations outside the U.S. and China are investing more in sovereign cloud IaaS to gain digital and technological independence," said Rene Buest, Sr Director Analyst at Gartner. "The goal is to keep wealth generation within their own borders to strengthen the local economy."
"Governments will remain the main buyers to meet digital sovereignty needs, followed by regulated industries and critical infrastructure organizations, such as energy and utilities and telecommunications," said Buest.

Europe is forecast to surpass North America in sovereign cloud IaaS spending in 2027.

Regionally, Middle East and Africa (89%), Mature Asia/Pacific (87%) and Europe (83%) are projected to record the highest growth in sovereign cloud IaaS spending in 2026. While China and North America are forecast to be No 1 and No 2. in spending in 2026 at $47 billion and $16 billion respectively, growth for both will be in the 20 percent range. Europe is forecast to surpass North America in sovereign cloud IaaS spending in 2027 (see Table 2).

Geopatriation to provoke cloud provider shift

Geopatriation is becoming a reality. Gartner estimates that due to an increased desire for geopatriation projects, sovereign cloud IaaS spending will shift 20% of current workloads from global to local cloud providers. In addition, 80% of the sovereign cloud IaaS spend will come from net new digital solutions or legacy workloads waiting to be migrated to a cloud environment.

Hyperscalers face mounting pressure as local cloud providers gain share and governments demand greater platform regionalization to meet regulatory and national security requirements. "To compete for local customers' cloud business, large cloud providers must seriously acknowledge the sovereignty concerns and requirements per country, and act accordingly. Solely treating digital sovereignty as a pure security, regulatory and compliance topic is not enough," said Buest.

Global Public Cloud Spending to Surpass $1 trillion in 2026

Global spending on public cloud services is forecast to surpass $1 trillion in 2026 -growing over 21% - and is expected to double by 2029, according to the latest update of International Data Corporation's (IDC) Worldwide Software and Public Cloud Services Spending Guide.

GROWTH IS DRIVEN by enterprise-wide application modernization, accelerating adoption of AI-enabled platforms, and rising demand for secure, scalable digital infrastructure across industries.

What is happening in the global public cloud market in 2026?

Growth of public cloud services is led by continued expansion of Software-as-a-Service (SaaS) and a sharp acceleration in Platform-as-a-Service (PaaS) adoption as organizations scale AI development, data platforms, and cloud-native application environments.

Global public cloud market at a glance

- Total public cloud services spending: >$1 trillion (+21% YoY)
- Largest deployment category: SaaS, more than half of total spend
- Fastest-growing deployment category: PaaS, +37% YoY
- Top spending industries: Banking, Software and Information Services, Retail (25% combined)
- Fastest-growing industries: Capital Markets, Banking, Software and Information Services
- Leading regions: United States ($647B), Western Europe ($255B), APeJC ($84B)
- Five-year CAGR above 20%: Middle East and Africa, Latin America, APeJC, Central and Eastern Europe, Western Europe

Why is public cloud spending accelerating?

SaaS Anchors Market Scale

SaaS will remain the largest public cloud deployment category in 2026, accounting for more than half of global spending. Looking at the largest secondary markets, investments are concentrated in IaaS, Enterprise Resource Management (ERM), and security software, reflecting enterprise priorities around cloud migration and cybersecurity, and core modernization.

PaaS fuels growth momentum

PaaS will be the fastest-growing deployment model, expanding over 37% year-over-year in 2026. Growth

is driven by rising demand for artificial intelligence (AI) platforms and application development software, as organizations adopt cloud-native environments to support generative and agentic AI, real-time analytics, and data-intensive workloads.

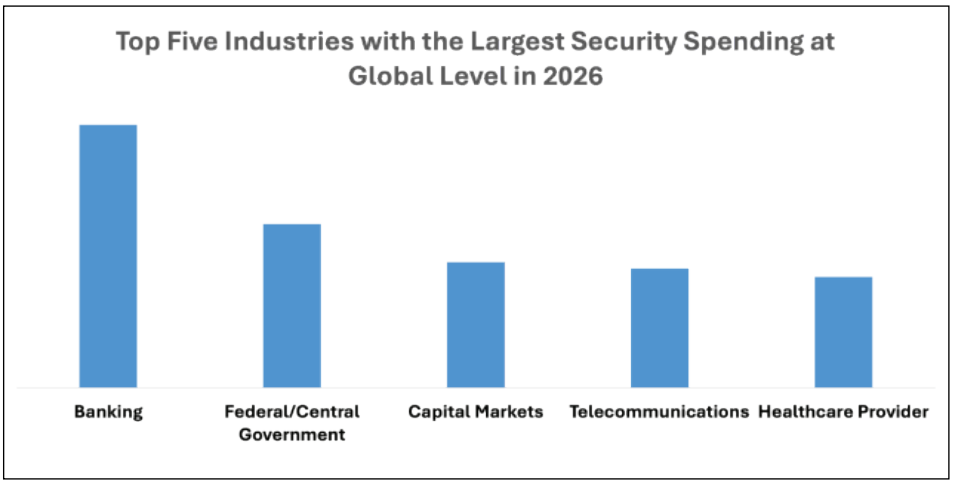

Industry Trends: Where is growth strongest?

Banking, software and information services, and retail will be the three largest public cloud–spending industries in 2026. The next five largest industries — Professional and Personal Services, Capital Markets, Media and Entertainment, Telecommunications, and Healthcare Providers — will account for more than a quarter of global spending.

Three fastest-growing industry for PaaS

- **Banking:** Institutions are modernizing core systems and deploying AI-driven risk, fraud, and real-time banking services.
- **Retail:** PaaS supports rapid development of applications for dynamic pricing, inventory optimization, and digital commerce.
- **Aerospace and Defense:** Increased defense budgets and geopolitical tensions are driving investment in secure cloud platforms supporting analytics, AI-enabled intelligence, and mission-critical systems.

Regional outlook

- **United States:** Spending will reach $647 billion in 2026, supported by large-scale enterprise migrations, hyperscaler AI infrastructure, investment, and strong demand from regulated industries.
- **Western Europe:** The market will total $255 billion, driven by cloud modernization programs, sovereign cloud adoption, and regulatory-led investment in data protection and AI governance.
- **APeJC:** Spending will reach $84 billion, fuelled by expanding digital economies, midmarket cloud adoption, and government-led digital transformation initiatives.

"Public cloud spending will continue to grow as cloud migration remains central for agility, resilience, and efficiency," said Andrea Minonne, research manager at IDC. "In aerospace and defense, higher budgets and rising geopolitical tensions are driving demand for secure, cloud-based platforms supporting advanced analytics and mission-critical systems, with increased defense spending across NATO members and escalating tensions in the Middle East accelerating investment in AI-enabled and security-focused cloud environments."

Global security spend to exceed $300 billion

Global security spending is projected to reach **$308 billion in 2026** and **$430 billion by 2029**, according to the latest forecast from the International Data Corporation's (IDC) Worldwide Security Spending Guide. The global security market is forecast to **grow 11.8% in 2026**, driven by increasing investments into unified, AI-driven security platforms and related services.

Global security market at a glance

- Total security spending: $308 billion (+11.8% YoY)
- Software: Largest tech group, >50% of spend
- Software: Fastest growth, +14% YoY
- Top industries: Banking, Federal/Central Government, Capital Markets
- Fastest-growing industries: Capital Markets, Media and Entertainment, Software and Information Services
- Leading regions: USA ($150 billion), Western Europe ($69 billion), APeJC ($26 billion)
- Fastest-growing regions: Middle East and Africa, Latin America, USA

What are the main technology drivers for security spending?

The largest technology group in 2026 is projected to be **Software**, accounting for more than half of the worldwide security spending. Identity and Access Management Software, Endpoint Security Software, and Security Analytics are expected to represent more than 50% of global Security Software spending this year. As threats—including those fueled by AI —become increasingly sophisticated, companies are prioritizing these tools to prevent breaches, protect critical assets, and gain actionable visibility across their environments.

Software is also forecast to be the fastest growing technology group in 2026, with an estimated year-over-year growth rate of 14% in 2026. Services are projected to follow closely, likewise expected to see double-digit growth this year. Cloud Native Application Protection Platform (CNAPP), Identity and Access Management Software, and Information

➤ IDC's Worldwide Security Spending Guide, February 2026

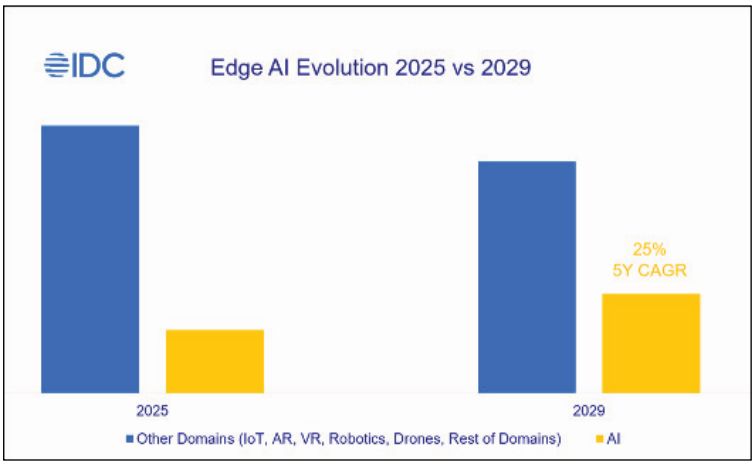

➤ Worldwide spending on edge computing reached $265 billion in 2025 and is expected to nearly double by 2029, according to IDC.

and Data Security Software will be the fastest growing Security Software technology categories. These technologies constitute the essential defense necessary to protect AI workloads, verify the identity of a non-human workforce, and ensure data safety in an era of AI-driven threats.

Among Services, **Managed Security Services** are expected to see the highest growth this year, allowing companies to help bridge the gap between escalating cyber-complexity and the persistent global shortage of in-house security talent.

"Organizations are moving beyond isolated security tools toward more integrated and intelligence-driven security architectures as threat complexity, regulatory pressure, and AI adoption accelerate," says Monika Soltysik, senior research analyst, Security & Trust at IDC. "Investment is increasingly focused on technologies that improve visibility, automate response, and strengthen identity and data protection across hybrid and cloud environments. Over the next several years, security strategies will increasingly prioritize operational resilience and platform consolidation as organizations seek measurable risk reduction rather than incremental tool expansion."

Key global dynamics

The global security market is shaped by AI-driven threats and increasing complexity of cyberattacks, prompting higher investments by companies in advanced security solutions. Geopolitical tensions and state-sponsored cyber operations are also intensifying risk and driving cross-border security spending. The United States will lead worldwide security spending in 2026, reaching $150 billion, driven by significant investments by the Financial Services, Healthcare, and Government industries. Western Europe will be the second largest market at $69 billion, pushed by intensifying regulatory and compliance requirements (e.g. NIS2, DORA, AI Act). APeJC will rank third with $26 billion, as rapid digital transformation and cloud adoption in the region will in turn trigger the required investments in security.

What are the key industry trends in security spending?

Banking, Federal/Central Government, Capital Markets, Telecommunications, and Healthcare Provider will be the top five industries in terms of global security spending in 2026, accounting together for more than one third of the total.

Capital Markets, Media and Entertainment, and Software and Information Services will be the fastest growing industries for security spending in 2026. Capital Markets will still be one of the primary targets for ransomware, fraud, and AI-driven cyberattacks – zero trust adoption, regulatory compliance automation, and AI-driven threat detection will be at the core of the security strategy of companies in this industry.

Media and Entertainment companies significantly rely on digital content distribution and cloud platforms and will increasingly focus their security investments on IP protection, piracy, and service disruption. As Software and Information Services companies are responsible for managing large-scale cloud infrastructure and client data as well for protecting the whole AI supply chain, a significant amount of spending will be dedicated to DevSecOps, CNAPP for multi-tenant environments, identity and access management solutions to secure users and services across platforms, and security analytics and automated incident response to handle large-scale threats.

Other relevant mentions include **Aerospace and Defense** and **High Tech and Electronics**. Given their high exposure to cyber espionage and nation-state threats, the two industries will continue increasing their investments in intellectual property and sensitive data protection throughout 2026. Supply chain security and third-party risk management will still be crucial for them to mitigate vulnerabilities across complex, global production ecosystems. Finally, the growing convergence of IT and OT environments will continue to drive demand for advanced protection of connected manufacturing systems and industrial infrastructure.

"The ongoing rise in cyberthreats and regulatory pressure will continue to drive global demand for resilient, sovereign, and compliant cybersecurity capabilities," says Stefano Perini, research manager, Market & Industries at IDC. "The strongest growth in security spending in 2026 is expected in the industries where protecting sensitive data, intellectual property, and critical infrastructure is most essential, and where the need for industry-specific security solutions is greatest. The growth will be greater for large companies, but it will also be significant for medium-sized and small enterprises, which are realizing that security is becoming an essential business enabler for them as well."

Edge computing global spending to grow at 15%

Worldwide spending on edge computing reached $265 billion in 2025, expected to nearly double by 2029, reflecting a pivotal expansion fueled by rapid Edge AI advancements that accelerate enterprise transformation and open new opportunities for service providers, according to the IDC Worldwide Edge Spending Guide.

"The combination of maturing edge architectures and rapid AI development is fundamentally redefining how organizations process and act on data," said Alexandra Rotaru, Data & Analytics Manager and WW Edge Spending Guide Product Lead at IDC. "We see enterprises and service providers shifting toward intelligent, distributed systems capable of realtime decisioning and automation at scale. Edge AI is no longer experimental, its impact is already visible across industrial automation, smart retail, connected vehicles, and next-generation healthcare."

Edge as the foundational layer enabling multiple technology domains

IDC segments edge spending across more than 1,000 named enterprise use cases spanning six domains—AI, IoT, AR, VR, drones, and robotics—highlighting the expanding role of edge infrastructure in enabling advanced intelligent workloads. Artificial Intelligence is one of the fastest-growing domains in the forecast, reflecting the increasing need to process data and run complex models directly at the edge. As organizations deploy more AI-driven applications that require low-latency inference, realtime context, and resilient distributed architectures, edge computing becomes a critical enabler of innovation and value creation across industries.

Which edge-related technology dominates the market?

Hardware remains the dominant investment category early in the forecast period, propelled by the rapid adoption of AI-accelerated infrastructure and increasingly sophisticated edge systems. This momentum reflects the growing need for realtime computing at the point of data creation, as organizations deploy heavier edge architectures capable of supporting advanced AI workloads.

However, the aggregate **Services** segments (including **Provisioned** and **Professional Services**) are expected to surpass hardware's share by 2029. Within Provisioned Services, Infrastructure as a Service (IaaS) remains the fastest-growing category as organizations increasingly rely on scalable, flexible, and cost-efficient consumption models to meet rising AI-driven compute requirements at the edge. As more workloads shift toward distributed and AI-enabled environments, asaservice offerings provide a critical foundation for managing capacity, optimizing costs, and accelerating deployment across diverse industry use cases.

Impact across industries

The **Retail & Services** sector remains the largest contributor to global edge spending, driven by rapid adoption of AI and IoT-related use cases such as video analytics and realtime operational optimization. **Manufacturing & Resources** follows, with AI-enabled quality control, predictive maintenance, and autonomous material handling increasingly dependent on realtime edge processing. **Financial services** is the fastest-growing sector, supported by AI-driven fraud analysis requiring low-latency, distributed architectures.

Service Providers also continue to scale investments in AI-ready edge platforms—including MEC, CDN

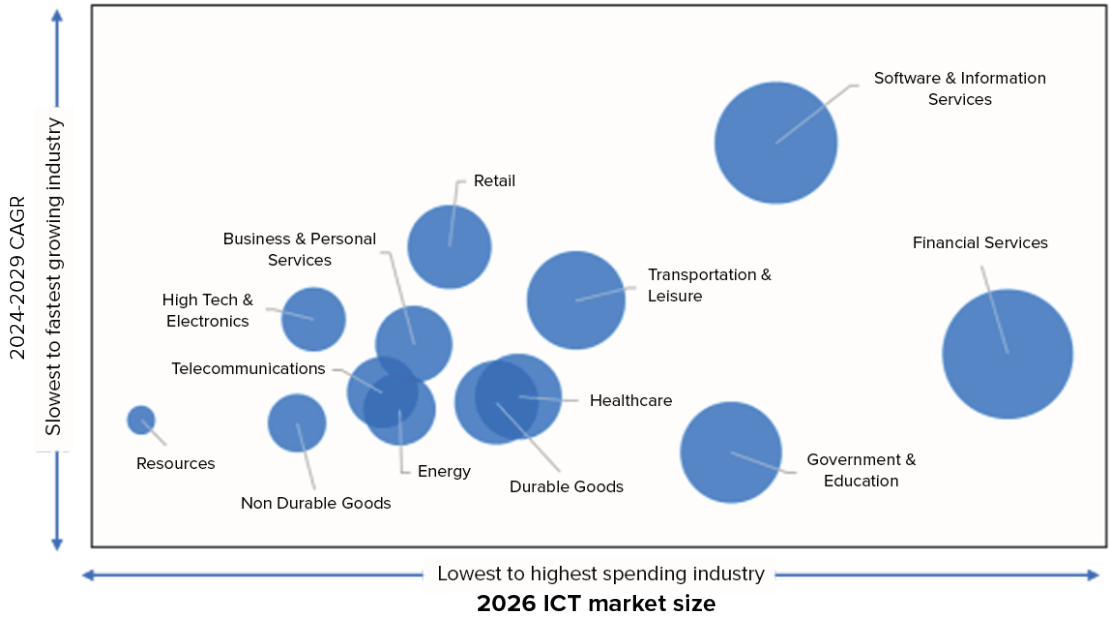

➤ Worldwide ICT spending market size and forecast by industry, 2024-2029

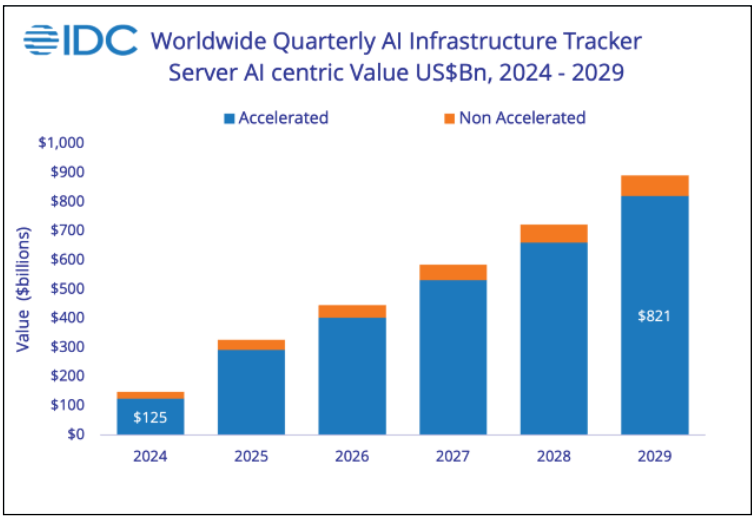

and virtualized network functions—which are on track to represent nearly one-third of the total market by 2029.

Regional development driven by the USA and Western Europe

North America will remain the top region for edge spending, driven by rapid adoption of EdgeAI workloads and advanced infrastructure deployments. Western Europe and China follow, supported by strong industrial and public-sector investment in AI-enabled applications. The United States and Latin America are expected to record the fastest growth, as organizations across both regions scale low-latency, AI-driven use cases that benefit from distributed edge architectures.

Global ICT spend to reach $4 trillion

Global spending on information and communications technology (ICT) is forecast to reach $4 trillion in 2026 and surpass $6 trillion by 2029, according to the latest update of the International Data Corporation's (IDC) Worldwide ICT Spending Guide Enterprise and SMB by Industry. Excluding the consumer segment, the ICT market is projected to grow 10% in 2026, driven by rapid adoption of Artificial Intelligence (AI) platforms, which will grow over 70% by the end of the year.

What are the main technology drivers for ICT spending?

Software emerges as the largest technology group in 2026, absorbing more than 33% of global ICT spending. This surge is primarily fueled by robust investments in enterprise resource management (ERM) applications, security software, and production and operations applications, which together will account for over one-third of total global software expenditure. These categories reflect enterprises' ongoing priorities around operational efficiency and cybersecurity resilience. Hardware is set to be the fastest growing technology group, with a projected year-over-year growth rate of 15% in 2026. This expansion is driven by significant momentum in non-x86 servers, the proliferation of wearables, and increased adoption of Infrastructure-as-a-Service (IaaS). This underscores the rising demand for specialized compute infrastructure, edge devices, and scalable cloud resources to support next-generation workloads and hybrid environments.

Global ICT market at a glance

- Total ICT spending: $4 trillion (+10% YoY)
- Software: Largest tech group, >33% of spend
- Hardware: Fastest growth, +15% YoY
- Top industries: Software and Information Services, Banking, Retail (>$1 trillion combined)
- Fastest-growing sectors: Software and Information Services, Media and Entertainment, Retail
- Leading regions: USA ($2 trillion), Western Europe ($908 billion), China ($355 billion)
- Five-year CAGR >10%: USA, China, Latin America

"We are entering a new phase of the AI-everywhere journey: the era of expectations and reckoning," said Andrea Siviero, senior director at IDC. "The excitement of experimentation is giving way to a sharper focus on accountability, value creation, and productivity impact. In 2026, enterprises will demand that AI and digital investments demonstrably improve processes, accelerate decision making, and ultimately drive business growth across the organization."

Key global dynamics

Major global events in 2025, including escalating trade tariffs, security threats, and the US government shutdowns, prompted organizations to accelerate investments in AI-driven optimization and security. Geopolitical tensions and supply chain disruptions accelerated digital transformation and security spending. The United States will lead global ICT spending in 2026, reaching $2 trillion, driven by its large enterprise base and rapid adoption of cloud and AI technologies. Western Europe will be the second largest market at $908 billion, supported by regulatory-driven modernization and accelerated AI adoption in sectors like banking and manufacturing. China will follow with $355 billion, fueled by government-led digital infrastructure and expansion of smart manufacturing and AI platforms.

What are the key industry trends in ICT spending?

Software and information services, banking, and retail will be the three largest spending industries and together will represent over $1 billion in ICT spending in 2026. The next five largest industries—federal central government, telecommunications, media and entertainment, healthcare provider, and high tech and electronics—will together account for more than 20% of worldwide spending.

"As companies continue to invest in automation, industries such as aerospace

and defense, insurance, and software and information services are poised to accelerate spending in AI platforms the fastest," said Andrea Minonne, research manager at IDC. "In aerospace and defense, escalating geopolitical cross-regional tensions and heightened security concerns are prompting governments to expand defense budgets".

Looking at other fast-growing vertical markets, insurance firms are leveraging AI platforms to manage rising claims volumes, improve risk modelling, and deliver real-time health insights, responding to both regulatory demands and evolving customer expectations. Software and information services are at the forefront of deploying generative and agentic AI to automate business processes, enable intelligent workflow orchestration, and support scalable digital transformation initiatives.

Meanwhile, high tech and electronics, automotive, and consumer goods are navigating ongoing frictions including trade tariffs and talent shortages and are prioritizing nearshoring strategies, factory automation, and advanced analytics. These investments are enabling organizations to mitigate supply chain risks, optimize production, and maintain competitiveness in a volatile global market.

AI infrastructure spending reached a record $86 billion

Worldwide spending on artificial intelligence (AI) infrastructure reached a record **$86 billion in Q3 2025**, marking a sustained investment cycle as platform providers scale capacity for training and inference workloads, according to the IDC Worldwide Quarterly Artificial Intelligence Infrastructure Tracker.

The record performance in Q3 2025 signals a shift from initial pilot phases into a multi-year expansion, **with full-year 2025 spending projected to reach $334 billion and more than $902 billion by 2029**. Growth is expected to remain above 30% annually through 2027, before moderating into the mid-20% range in the latter years of the forecast.

> The United States is expected to remain the largest AI infrastructure market in 2025, accounting for approximately 76% of global spending with investment projected to grow from $254 billion in 2025 to nearly $708 billion by 2029.

The results highlight the central role of accelerated compute as enterprises and cloud providers move to support increasingly complex AI workloads.

Why AI infrastructure grew in Q3 2025

Growth in Q3 was driven by massive capital investment from AI platform providers and hyperscalers, with hardware categories expanding to meet the rigorous data demands of large language models (LLMs).

- Servers dominated the market, accounting for $84.0 billion (nearly 98%) of total AI-centric spending.
- Cloud and Shared Deployments represented over 86% of the market, reflecting the industry's reliance on scalable infrastructure.
- AI-Centric Storage reached $1.76 billion, driven by the need for high-performance repositories for model training and inference checkpoints.

Accelerated compute and GPU demand accelerate growth

Accelerated servers remain the foundation of AI infrastructure. While YoY growth moderated from earlier peaks in 2025, spending levels remained elevated, signaling sustained demand rather than a pullback in investment. Both x86 and non-x86 systems saw strong adoption as AI platforms scaled infrastructure to meet rising training and inference requirements.

Regional growth led by the united states and china

Geography played a major role in the Q3 surge, with two regions leading the global investment landscape:

- The United States is expected to remain the largest AI infrastructure market in 2025, accounting for approximately 76% of global spending with investment projected to grow from $254 billion in 2025 to nearly $708 billion by 2029. Growth is driven by hyperscalers and AI platform leaders expanding large-scale data center capacity.
- China (PRC) is expected to be the second fastest-growing major region over the forecast period, with spending projected to increase from $39.1 billion in 2025 to more than $139 billion by 2029, supported by continued investment in domestic AI platforms and sovereign AI initiatives.

"The AI infrastructure market has clearly moved beyond an initial deployment phase into a sustained expansion cycle," said Lidice Fernandez, group vice president, Worldwide Enterprise Infrastructure Trackers at IDC. "Public investment signals from leading AI platform providers point to multi-year commitments to infrastructure expansion, particularly around accelerated compute. These investments reflect long-term confidence in the growth of AI workloads across consumer, enterprise, and research use cases."

From ‘automated’ to orchestrated: Closing the real productivity gap in the enterprise

Over the past two decades, enterprises have poured billions into automating business operations. Payroll runs itself. IT accounts are provisioned in seconds. Yet employees are still losing up to seven hours a week to tasks like chasing approvals, booking travel or filing expenses. That disconnect reveals a critical oversight; most digital transformation has automated around the employee, not for them.

BY ROBIN SMITH, CTO OF PERK

THIS IS THE REALITY of shadow work - the everyday admin that quietly drains time, morale and momentum. It’s costing UK businesses £95 billion in lost productivity each year. And because it rarely shows up in dashboards or KPIs, it continues unchecked, even as businesses ramp up automation spending.

If the first wave of enterprise automation was about streamlining systems, the next must focus on orchestrating work - designing seamless, AI-driven experiences across the tasks that make up an employee’s day.

The automation gap is bigger than it looks

Businesses haven’t ignored automation; they’ve just deployed it in the wrong places. Core systems are optimised but the tasks that sit between them - managing expenses, coordinating travel, chasing invoices - remain stubbornly manual.

This misalignment has created a widening automation gap; a mismatch between where automation has been applied and where it would have the greatest impact on productivity and employee experience. While systems run efficiently in the background, employees juggle fragmented workflows. According to our research, they use four different tools on average to manage shadow work, yet just 7% say those tools are well integrated.

Because shadow work cuts across departments, it often escapes ownership. It’s too operational for IT, too invisible for HR, and too tactical for the C-suite. As a result, businesses invest heavily in automation while neglecting the points of greatest friction.

Why rationalising the stack is a leadership decision

The natural response to complexity has often been to add more tools. A new app for travel. A separate platform for expenses. Another layer for approvals. But more tools don’t always mean more productivity - they often introduce more logins, learning curves and policy confusion.

Instead of solving the problem, fragmented tech stacks shift the burden onto employees, who spend their time stitching systems together instead of doing the work they were hired for.

With 67% of global leaders already committed to investment in automation, it’s time to refocus. Leaders must stop layering tools and start embedding intelligence into the operational core, where workflows, policies and systems converge.

In fact, our report shows that the companies seeing the biggest gains from automation are those that apply it where the productivity drag is greatest - tasks employees say are frequent, frustrating and low-value. These are not abstract problems, they’re recurring bottlenecks that show up every week in expense reports, travel bookings and team coordination. They have a known cost and a measurable upside when removed.

Rationalising the stack is about focus rather than chasing the latest platform. The businesses that succeed will be

the ones that prioritise simplicity over scale, and employee experience over software sprawl.

Orchestration: Automation that works for people

Fixing the automation gap isn't about adding more automation, it's about aligning the tech and problem you want to solve. And this may not mean adding another system. It starts with looking at how work is really done, and how tasks move across finance, operations, HR and beyond.

Shadow work often exists in the grey areas between these functions, where no system is fully responsible. Orchestration bridges these gaps. Instead of forcing employees to jump between apps or chase information, AI-native orchestration coordinates tasks behind the scenes, enforcing policies and routing workflows automatically.

This allows employees to be quicker but also gives them more clarity. People stop losing time to distractions and start focusing on work that matters. We found that it takes 11 minutes to refocus after completing a shadow task. Multiply that across an organisation, and the case for orchestration becomes clear.

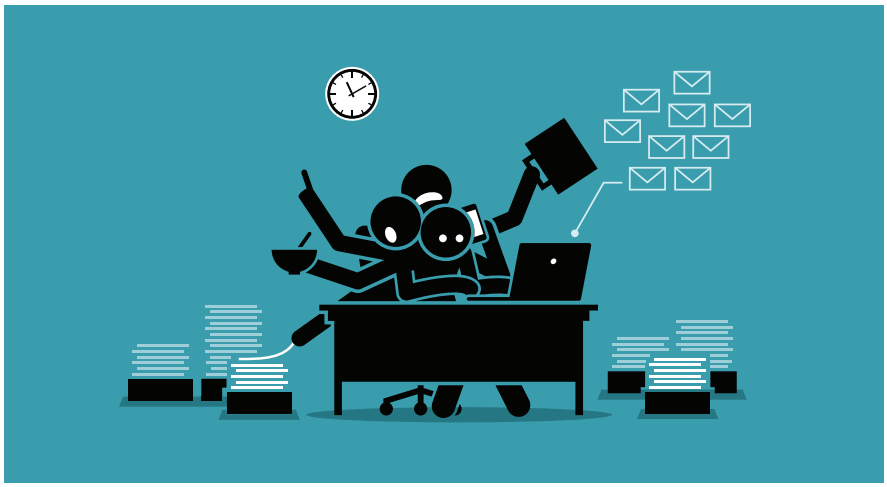

Shifting the focus

Most companies aren't underpowered, but rather they're misaligned. Automation has happened, but in ways that haven't freed employees to do their best work.

To unlock real productivity, leaders must rethink how work is designed. That means simplifying systems, orchestrating experiences across departments and freeing people from the hidden admin that slows them down.

The tools are ready and the core problem is clear. What's needed now is a shift in focus, from building smarter systems to creating smarter workflows that power real work.

How to build an AI SRE agent that solves production incidents like a team of engineers

The next generation of AI SRE agents must investigate like engineers, forming hypotheses and tracing causality. Grounded in real incidents and telemetry, these systems can evolve from noisy tools into trusted operational partners.

BY PEJMAN TABASSOMI, FIELD CTO FOR EMEA AT DATADOG

SITE RELIABILITY ENGINEERS (SREs) are constantly flooded with alerts in large-scale, distributed environments, where every service, API, and infrastructure layer can fail. As outages span multiple services and generate increasingly complex telemetry, pinpointing root causes has become significantly more difficult.

To address this, many enterprises have turned to AI agents to automate incident management, aiming to reduce alert fatigue and shorten mean time to resolution. In practice, however, these systems often struggle to cut through the noise. They tend to summarise alerts rather than properly identifying the root cause, leaving engineers to do the real investigative work.

The next step is to build AI SRE agents that think more like human engineers: capable of forming hypotheses, testing them against telemetry, and following evidence to uncover true root causes. Achieving this requires benchmarking against real incidents, ensuring agents can adapt to the complexity and unpredictability of production environments. Only then can AI move from a simple automation tool to a trusted teammate in incident response.

From summarisation to investigation

Traditional monitoring systems excel at detecting incidents but often fail to provide meaningful insight. When multiple alerts trigger simultaneously, the key question is not just what happened but why? AI SRE agents should replicate the investigative workflow of human engineers: forming hypotheses, testing them, and refining their understanding until the underlying cause is clear.

Modern observability platforms make this possible. Agentic AI can ingest telemetry data across various services, investigate incidents, and recommend or even implement remediation steps. By capturing each stage of reasoning (from task planning and tool invocation to data retrieval and decision-making) organisations can transform AI from a black box into a transparent, accountable participant in incident management.

Ground agents in real incidents (not synthetic tests)

A common mistake in AI operations tooling is evaluating agents only against simplified or synthetic scenarios. Real-world systems fail unpredictably: degradations may occur gradually, cascading failures can emerge, and noisy metrics often obscure the true root causes.

Benchmarking against actual production incidents is therefore essential. Observing how AI responds to real outages allows teams to refine reasoning paths, reduce false positives, and strengthen root-cause analysis. This creates a continuous improvement loop, where each investigation

improves the agent's ability to tackle future incidents.

Access to large-scale, real-world telemetry provides another advantage. Real-world data from production environments enables AI SRE agents to recognise meaningful patterns, trace causal relationships and filter out irrelevant signals (capabilities that are difficult to develop using hypothetical scenarios alone).

Following causality, not noise

Distinguishing correlation from occurence is a key skill for SREs, and the same principle applies for AI agents. Alerts and anomalies rarely exist in isolation; they propagate through dependencies across services and infrastructure.

Effective AI must trace these signals along their relationships to uncover root causes, rather than treating each alert as a separate issue.

Agents that can correlate telemetry across infrastructure, service chains, and applications gain a holistic view of incidents. This enables them to filter out irrelevant signals, focus on meaningful causal relationships, and carry out more insightful investigations.

Causality-driven reasoning depends on access to high-quality, real-world production telemetry.

Observability must extend beyond uptime and latency to include metrics such as model accuracy, data integrity, and operational behaviour. With this broader context, AI can detect complex failure modes (including those introduced by AI-driven workloads) without being misled by superficial signals.

As agents learn from real incidents and adapt to changing conditions, they evolve from reactive responders into proactive problem solvers. The result is a self-reinforcing cycle in which exposure to production data sharpens reasoning, speeds investigations, and improves resilience at scale.

From tool to teammate

The ultimate goal is not to replace engineers but to amplify their expertise. Well-designed agents can streamline investigations, highlight likely root causes, and confidently recommend (or even execute) remediation steps.

When grounded in real telemetry, benchmarked against actual incidents, and designed to prioritise causality over noise, AI becomes more than a monitoring tool. It becomes a trusted teammate that investigates, reasons, and learns alongside the engineers it supports.

In high-scale environments, minutes of diagnostic delay translate directly into revenue loss and customer impact. As digital estates grow, organisations that treat AI as an investigative partner rather than a passive summarisation tool will excel in operational resilience.

Distinguishing correlation from occurence is a key skill for SREs, and the same principle applies for AI agents. Alerts and anomalies rarely exist in isolation; they propagate through dependencies across services and infrastructure. Effective AI must trace these signals along their relationships to uncover root causes, rather than treating each alert as a separate issue.

By combining telemetry, experimentation, and structured reasoning, these organisations can move from reactive firefighting to proactive incident prevention. The result is a future where outages are minimised, reliability is maximised, and engineers can focus on strategic challenges rather than firefighting alerts.

Unlocking the true business value of modern log management

Recent outage highlights just how essential it is for businesses to have clear, real-time visibility into what's happening inside their digital systems. Logs, defined as the record of activities and operations that occur within a computer system have become one of the most critical sources of truth for understanding system activity, from server actions and user interactions to error messages and early signs of a cyberattack.

BY MALA PILLUTLA, VICE PRESIDENT OF SALES FOR LOG MANAGEMENT, DYNATRACE

LOG MANAGEMENT IS, therefore, about turning raw system logs into actionable insights. As digital systems grow in scale and complexity, logs have evolved from a backroom tool into a critical driver of reliability, performance and security across an entire business.

Research has found that the majority (87%) of organisations claim to use logs as part of their observability solutions. However, while log usage is near-universal, a question remains over whether businesses are unlocking their full value. Collecting logs is one thing but interpreting them is another.

Where traditional log management falls short

As business digital estates grow more complex, the volume of logs generated across applications, infrastructure and business services has exploded. However, more logs do not automatically mean more insight. In fact, many teams are overwhelmed by sheer volume, struggling to separate meaningful signals from background noise. This overload creates noise that makes it difficult to identify urgent issues, leaving IT and Security teams on the back foot during critical incidents and proactive response.

The problem is as much about cost as complexity. Storing and managing log telemetry without a clear purpose often leads to escalating expenses that outpace the value delivered. Traditional licensing and infrastructure models add to the problem. They often make log management feel like a financial liability than a strategic advantage.

Another common constraint is fragmentation. Logs often live across multiple tools, with different interfaces and storage models, slowing root cause analysis and complicating cross-team collaboration. In a cloud-native world

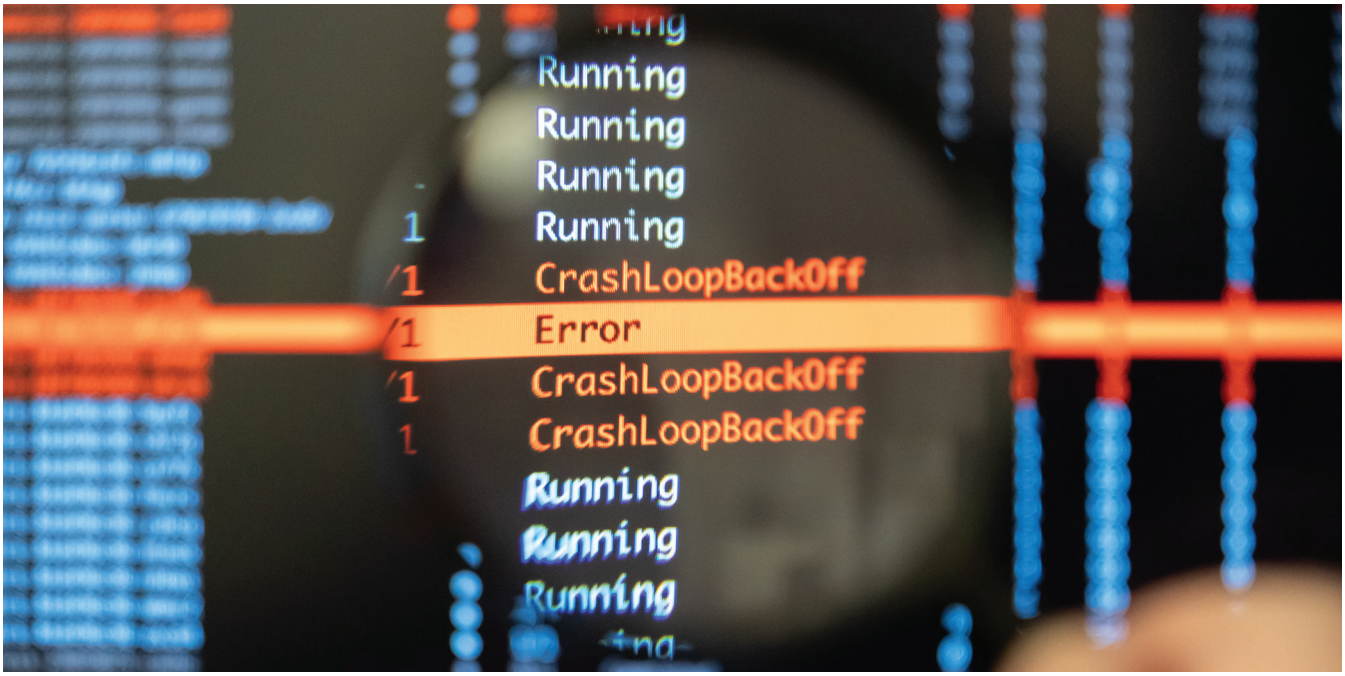

Forward-thinking organizations decide which logs to capture, which to discard, and how to route them most effectively. This helps control costs and ensures that attention is focused on the telemetry that delivers the greatest value. When organizations find this balance, log management evolves from a tactical task to a strategic capability that strengthens both performance and resilience.

where speed and scale are vital, this siloed approach is out of step with modern business needs.

Together, these shortcomings point to the need for a smarter approach—one that focuses on clarity, efficiency, and value.

From raw data to meaningful insight

Taking a smarter approach to log management starts with a shift in perspective. Rather than treating logs as an endless stream of technical data, leading organizations use them as a lens to understand how their digital ecosystems truly perform. The real value lies in not collecting everything but in knowing what matters and identifying which logs drive resilience, security, customer experience, or compliance, and filtering out the rest.

AI is becoming an essential part of this process. Modern techniques can detect anomalies, trace issues back to their root cause, and even trigger automated fixes. This reduces manual investigation and accelerates recovery, allowing teams to move from firefighting to foresight.

Equally important is being selective. Forward-thinking organizations decide which logs to capture, which to discard, and how to route them most effectively. This helps control costs and ensures that attention is focused on the telemetry that delivers the greatest value. When organizations find this balance, log management evolves from a tactical task to a strategic capability that strengthens both performance and resilience.

Log intelligence in the context of observability

Log intelligence on its own is valuable, but it is only part of the story. The next frontier is AI powered observability, uniting logs with metrics that track performance, traces that map interactions, and events that reveal key system changes. Combined in a single platform, these data types give teams a complete picture - connecting technical performance with genuine business impact and moving from a view of what happened to an understanding of why it happened and how to respond quickly.

Consider a global telecommunications provider that recently re-evaluated its log strategy. Managing more than 15TB of logs every day, stored for long periods and spread across thousands of dashboards, the team was buried in dashboards and redundant data. By consolidating logs within a broader observability framework and replacing static alerts with intelligent detection, they cut through the noise across its systems. Able to focus on the signals that mattered most, the organization improved uptime, speed, and overall resilience.

This example shows that observability delivers its greatest value when it helps teams cut through complexity. With logs feeding into a single platform, data becomes easier to interpret and act on, transforming technical insight into business intelligence.

Realising the full potential of log management

Modern log management gives organizations the context they need to turn massive volumes of data into meaningful insight. Organizations that harness AI, automation, and broader observability, gain a clearer view of how their technology is supporting their goals. Enterprises can analyse faster, automate smarter, and innovate with confidence.

True modernization comes from changing how teams think about data. Now is the time to review current strategies, identify gaps, and adopt modern platforms that integrate AI, context, correlation, and smarter telemetry management practices because organizations can no longer afford to treat log management as a background IT task. The companies that thrive will be those that treat logs not as exhaust from their systems, but as evidence of how their business thinks and performs. By bringing intelligence to the data they already have, they will turn observability into a source of continuous advantage and understand their business like never before.

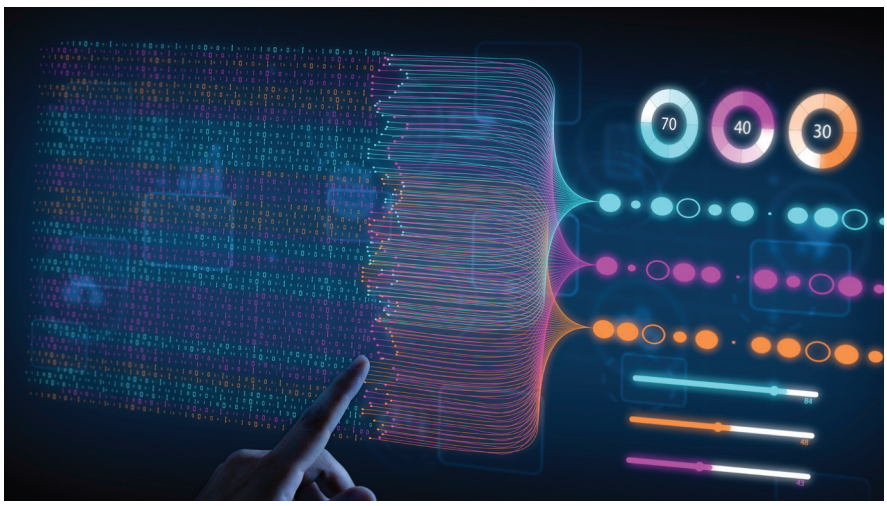

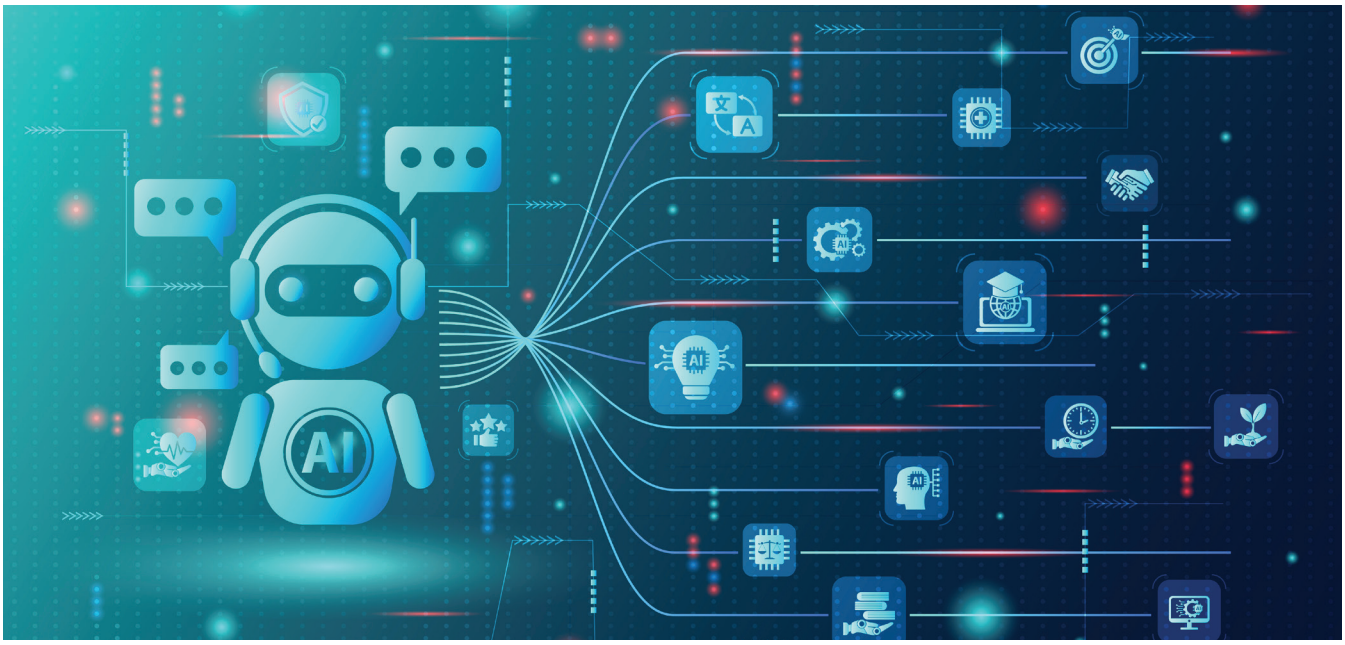

Why agentic AI's next challenge is making systems work together

Agentic AI will continue to advance, that's a given. But its trajectory and benefit inside the enterprise will depend on how well organisations adapt.

BY KARTHIK SJ, GENERAL MANAGER OF AI AT LOGICMONITOR

ASK ANY IT team how they usually spend their week and you'll likely get the same answer – firefighting.

It's not the urgent, high-stakes level of putting out fires, but the slow-burn of everyday operational fires of triaging an influx of alerts, tracking down duplicate tickets and moving information from A to B because they can't talk to each other. The work inevitably piles up, creating a constant heat that pulls teams away from the strategic work - and potentially larger fires to put out.

Agentic AI is already changing that, helping teams handle the repetitive operational work that takes up the bulk of their daily bandwidth, spotting problems early and figuring out what they mean in context before acting on them without the usual manual oversight needed.

But we need to recognise that this is the easy part. The harder question is what happens when teams want to move beyond one agent doing one thing to expand across different workflows and business functions.

The issue isn't whether any single agent can deliver. It's whether multiple agents, built by potentially different departments or vendors, can operate coherently together to plan, act and embed automation directly into how an organisation runs.

As agentic adoption climbs, so does the complexity

The shift towards agentic systems is no longer confined to IT. Increasingly, functions including finance, customer support and operations, sales, marketing, risk and compliance are seeing these systems shape the order in which work happens and the decisions that follow.

Investment patterns for agentic systems in the US reinforce this, with leaders carving out budgets (around 25%) dedicated for agentic capabilities rather than treating them as a secondary investment or sub-feature of 'general AI'. The same direction of travel is visible in the UK, with AI adoption moving into the mainstream for a sizable portion of businesses, with over a quarter reporting they were already using some form of AI in 2025.

As that baseline becomes normal, the focus shifts naturally towards the next layer of maturity: how agent-driven systems can operate reliably in environments that are already complex.

Scaling steadily and navigating the obstacles

This is where the real test begins for IT leaders. What works well in a small, neatly defined deployment can start to buckle once agents are asked to coordinate across platforms and teams.

In other words, the question is no longer whether an agent can complete a task, it's whether a growing number of agents can make compatible decisions when they share responsibility for the same outcomes.

The problem is rarely that agents fail outright. It is that their decisions begin to overlap and override one another. Each may act sensibly based on its local context but without a shared view, the combined outcome can be messy, creating duplication, contradictions and new issues that teams have to spend even more time untangling.

Frictions created by these conflicting processes are not a result of agents being unable to perform, its because they weren't given the right guardrails on how to connect and work off one another.

Interoperability should be seen as the foundation

As agentic systems are plugged into core platforms such as CRM, ERP and service operations, those knock-on effects become unavoidable.

Actions taken in one area ripple into another. At that point, simply adding more intelligence does not improve performance and can even make it worse by increasing the volume of decisions without improving alignment.

Agentic systems don't behave like traditional software, where the same input reliably produces the same output. They take in new information, adjust what matters most, and change course as conditions shift. That's the point, but it also makes it harder to keep multiple systems aligned.

When scaling, progress depends on shared context. Agents need enough awareness of one another's actions and intent to avoid working at cross purposes. Without it, activity can increase without much material improvements.

The organisations that are progressing at pace, bake interoperability in from day one. They decide up front, which agents are allowed to take action and how information moves between systems instead of hoping it sorts itself out over time.

As agents are trusted with more critical workflows, end-to-end visibility stops being a nice-to-have.

What does this mean for IT leaders?

IT leaders will see the strongest ROI when agentic systems are allowed to work across platforms rather than being boxed into a single workflow; value shows up when agents can connect data and decisions end-to-end so the organisation can respond as one system instead of a set of disconnected teams.

That puts integration and visibility at the top of the agenda. Research repeatedly points to the same constraints: fragmented environments and capability gaps that make it difficult to take AI beyond pilots. As agentic adoption matures, interoperability becomes less of a technical preference and more of an operating requirement.

For executives, the difference between experimentation and impact is measurable and will be clearly seen. Faster incident resolution and real productivity gains make it obvious where these systems are improving performance and where they are introducing new complexity that still needs tightening.

Agentic AI will continue to advance, that's a given. But its trajectory and benefit inside the enterprise will depend on how well organisations adapt. The next wave of impact won't come from smarter individual agents, it will come from smarter ways of getting those agents to work together efficiently.

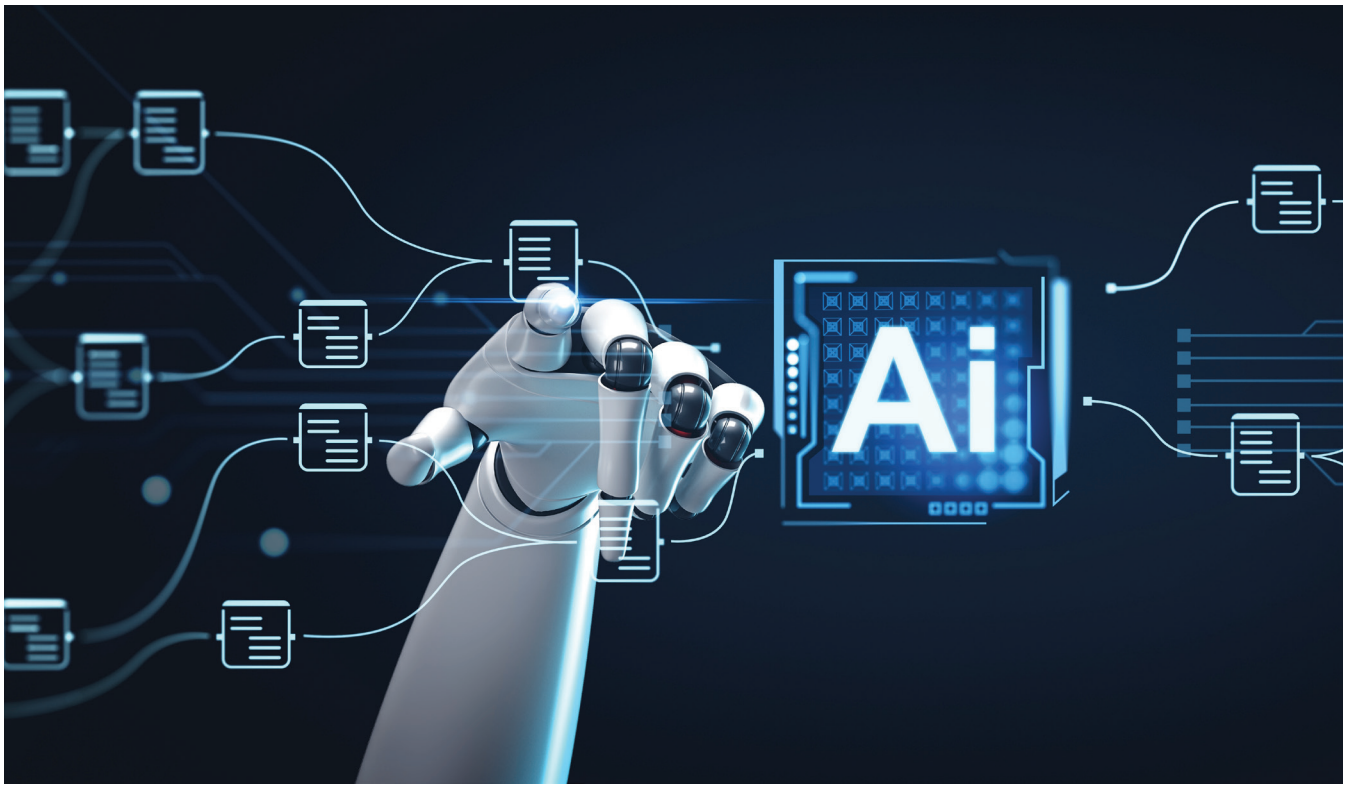

What's observability? And why should I care if I've got AI?

Technology was supposed to make everything easier. Faster decisions, smarter systems, leaner operations. But for many leaders, the reality looks very different: rising costs, swelling cyber risk, and a tangle of legacy and multi-cloud complexity that's only getting harder to manage. AI is now promising to help solve this — but in truth, AI has its work cut out.

BY SAMMY ZOGHLAMI, SVP EMEA AT NUTANIX

EVEN THE MOST advanced projects are feeling the strain. Gartner predicts that over 40% of agentic AI projects will be cancelled by the end of 2027, derailed by spiralling costs, poor ROI, or inadequate risk controls. It's not that AI cannot deliver value, it's that the foundations beneath it were not built for what is coming.

The truth is, much of the frustration around AI isn't really about AI at all. It's about those underlying systems and an organisation's ability to actually see what is happening across those systems. For all the hype around intelligent agents and autonomous workflows, success still depends on something far less glamorous; the performance, visibility, and resilience of the platforms those models run on. When infrastructure cannot keep pace, costs rise, performance dips, and complexity multiplies. That is where the cracks begin to show.

Three recurring pain points keep surfacing.

AI workloads are starving for data

Modern models devour data. Training runs and retrieval-augmented generation (RAG) pipelines depend on high-throughput access to files, objects, and vector data spread across hybrid environments. Yet traditional storage systems were not built for this pace. I/O bottlenecks throttle performance, GPUs sit idle waiting for data, and every wasted second becomes wasted compute. For many enterprises, storage has become a tax on AI progress.

Observability isn't observability if it stops at infrastructure metrics

Most organisations can see CPU load, disk IOPS, and network latency but that's only half the picture. True observability means correlating those infrastructure signals with model behaviour - accuracy, drift, throughput, error rates, even cost per inference.

But performance on its own isn't enough. Without visibility, even the best-engineered systems are flying blind. Observability needs to go beyond dashboards and alerts. It has to connect the dots between infrastructure health and model behaviour. It is the ability to see how a GPU spike in one region affects inference latency elsewhere, or how network congestion is degrading model accuracy. When you can see everything, data, compute, and model performance, you can tune it, fix it, and ultimately trust it.

When data, compute, and models are scattered across clouds, this end-to-end view disappears. Teams end up reacting to symptoms, such as slower queries and rising bills, without understanding root causes. Observability, in simple terms, should answer one question. What's happening, why, and what should we do about it?

Fragility is a hidden threat

AI workloads are notoriously unforgiving. A single node failure, power fluctuation, or regional outage can derail production workflows, interrupt inference pipelines, and erode business confidence. Many enterprises still rely on manual failovers or untested disaster-recovery plans. True resilience means cross-region redundancy, automated recovery, and continuous validation because in AI, uptime equals trust.

These three issues are driving the cancellations, overruns, and disappointments Gartner warns about. And they're why performance and resilience, the two least glamorous parts of the stack, have suddenly become the most strategic.

So, what does good look like?

It starts with recognising that performance is a by-product of smarter architecture (and not necessarily better hardware). The best AI systems are fed by storage that can keep up. That means fast, scalable, and intelligent enough to balance cost with speed. When training workloads or retrieval-augmented generation pipelines hit the accelerator, data needs to move just as quickly. AI-optimised, tiered storage architectures do exactly that, feeding GPUs at line speed while still providing the durability and auditability needed for compliance.

But performance on its own isn't enough. Without visibility, even the best-engineered systems are flying blind. Observability needs to go beyond dashboards and alerts. It has to connect the dots between infrastructure health and model behaviour. It is the ability to see how a GPU spike in one region affects inference latency elsewhere, or how network congestion is degrading model accuracy. When you can see everything, data, compute, and model performance, you can tune it, fix it, and ultimately trust it.

And then there's resilience, the quiet hero of AI scale. The more distributed AI becomes, the more fragile it gets. Models are trained across regions, data flows across clouds, and a single outage can ripple through everything. The answer isn't redundancy for redundancy's sake, it is resilience by design. Dynamic workload migration, self-healing infrastructure, and continuous validation of failover processes. That's what turns AI from an experimental tool into an operational asset.

In truth, performance, observability, and resilience are inseparable. Without one, the other falters. Together, they define how prepared an organisation really is for AI at scale, not just for the pilot phase, but for the everyday reality of running critical workloads in production.

AI success depends on seeing your infrastructure as part of the intelligence. Leaders should start by asking tough questions about visibility and control. Can your teams trace data flows across every cloud? Do you know, in real time, how infrastructure decisions affect model performance? And are your recovery processes tested for when (not if) something fails?

The answers shape competitive advantage. The organisations that treat infrastructure as a living system, continuously tuned, instrumented and stress-tested, will be the ones that turn AI into a reliable engine of productivity. Because the future of AI isn't just about creating smarter models, it's about value. Otherwise, what's the point? And without smarter systems, there really is no point.

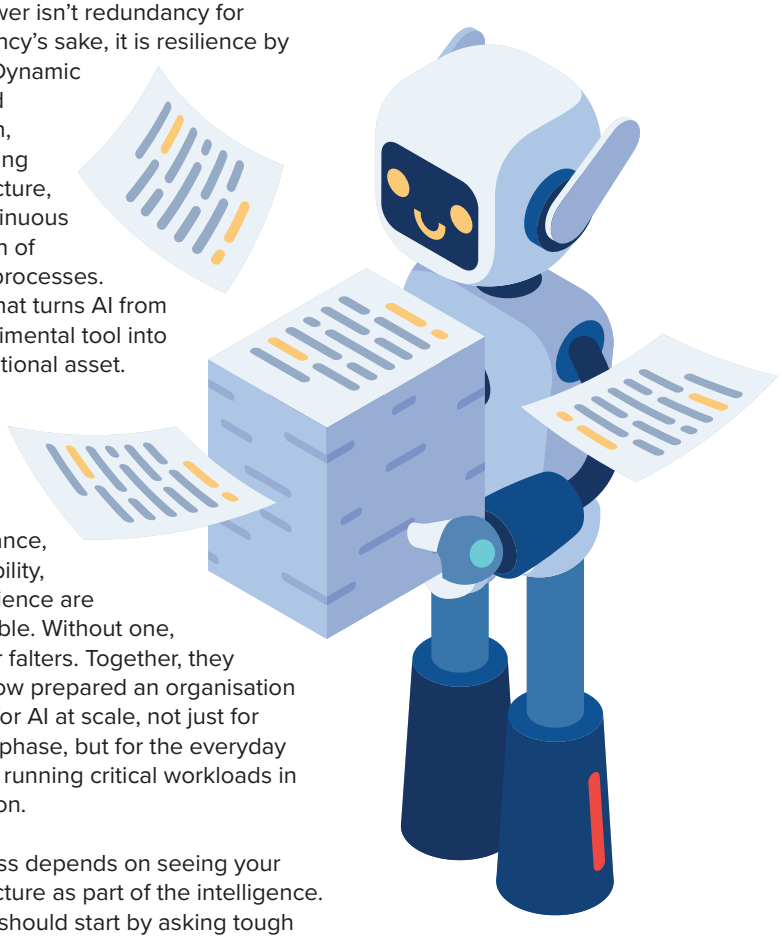

IT readiness in 2026 will be defined by data resilience - not digital ambition

For too long, organisations have mistaken digital acceleration for digital readiness. The rush to adopt AI, expand cloud usage and deploy new platforms has created an illusion of progress, but without resilient data foundations, much of that innovation remains fragile.

BY PAUL SPECIALE, CHIEF MARKETING OFFICER AT SCALITY

AS 2026 progresses, the defining question for IT leaders is no longer how fast they can transform, but how well their infrastructure can withstand disruption. The next phase of digitalisation will be shaped less by new tools and more by the durability of the data architectures that support them.

Data resilience means keeping data available, trustworthy and recoverable across failures, attacks and regulatory change. In practical terms, it shows up in a few simple tests: can you locate your most critical datasets, prove who can access them, and restore a clean copy within the time the business can tolerate?

From innovation to accountability

Boards are asking different questions: not only what is being deployed, but whether the underlying data environment can survive failure and return services to a known-good state. Innovation is no longer judged purely by speed; it is measured by operational resilience. That is pushing organisations to revisit the architecture beneath critical services, where reliability, governance and recoverability now sit alongside scalability.

AI readiness begins with data readiness

AI has accelerated this reassessment. Many enterprises initially approached AI as a compute problem, investing in GPUs and models. The stubborn bottlenecks have appeared in data pipelines - inconsistent metadata, fragmented storage environments and limited visibility into lineage and unclear ownership of datasets.

Without trustworthy data, AI outcomes are difficult to validate, reproduce or explain. That risk is now commercial as well as compliance-related. Stakeholders increasingly expect organisations to justify automated decisions and demonstrate provenance. The result is a renewed focus on governed pipelines, auditability and controlled access, recognising that AI success depends as much on data quality and availability as it does on algorithms.

Cyber resilience moves deeper into infrastructure

At the same time, cyber threats are evolving in ways that expose weaknesses in traditional architectures. Attackers are targeting backup systems and recovery processes themselves, forcing organisations to rethink how resilience is implemented. When recovery is compromised, “having backups” is not the same as being resilient.

Leading enterprises are building deeper architectural safeguards - immutable storage to prevent tampering, isolated recovery environments and regular/automated testing that validates restoration processes. Resilience is no longer defined by whether data is stored safely; it is defined by how quickly operations can return to a trusted state after disruption.

This change signals a broader shift in mindset. Cybersecurity is moving beyond prevention toward operational continuity, where infrastructure design plays as critical a role as security policy.

Sovereignty becomes an engineering priority

Data sovereignty is also moving from policy discussion to engineering

reality. Across Europe and other regulated regions, organisations must demonstrate clear control over where data resides and how it moves between jurisdictions. Compliance now requires more than geographic hosting; it demands verifiable governance embedded into infrastructure itself.

Hybrid and multi-cloud strategies are being reshaped by these requirements. IT teams must design environments capable of supporting regional controls without sacrificing flexibility or performance. The challenge is balancing compliance with agility, ensuring that regulatory demands do not slow innovation.

Rethinking cloud dependence

Cloud remains central to modern IT, but recent disruptions have encouraged organisations to take an honest look at concentration risk. Rather than abandoning cloud, enterprises are adopting more balanced strategies that prioritise portability and independence.

Architectures that support portability - through standard interfaces, consistent data services and clear exit paths - give organisations options when cost models change, performance requirements shift or sovereignty constraints tighten. Choice becomes a resilience feature.

Disaggregation and adaptability

Readiness is also being shaped by disaggregated infrastructure. Traditional systems designed around fixed workloads struggle to support the unpredictable growth patterns created by AI and analytics. Separating compute, capacity and metadata scaling allows organisations to respond more dynamically as demands evolve.

This architectural shift reflects a broader recognition that digital growth rarely follows a predictable path. Rather than building systems optimised for a single outcome, IT leaders are designing platforms capable of adapting continuously; reducing bottlenecks while maintaining performance.

The hidden risks organisations still overlook

Even with growing awareness of resilience, a few gaps continue to undermine readiness:

- Vendor lock-in that limits data mobility and complicates sovereignty and cost control.
- Recovery plans that are rarely tested end-to-end, leaving true recovery time unknown.
- Weak lineage and provenance, turning AI success into an audit and compliance risk.
- Data sprawl from rapid experimentation, expanding the attack surface and making governance harder.

Storage steps into a strategic role

Taken together, these pressures are elevating the role of storage from “plumbing toward strategic maturity”. The data layer is where security, governance and performance converge. Infrastructure decisions made today will determine how confidently organisations can adopt emerging technologies tomorrow.

This is not a call to slow down innovation. It is a shift towards maturity: resilient foundations enable faster progress because they reduce uncertainty, shorten recovery and make compliance demonstrable.

From digital ambition to digital durability

In 2026, IT readiness will be defined less by the pace of adoption and more by the strength of underlying infrastructure. Governed data pipelines, deeper cyber resilience, operational sovereignty and cloud independence are becoming baseline expectations rather than aspirational goals.

The organisations that pull ahead will be those that treat resilience as an advantage, and design for disruption, not assuming it will go away. In an era shaped by AI disruption and regulatory complexity, digital ambition alone is no longer enough. Digital durability is what will set tomorrow’s leaders apart.

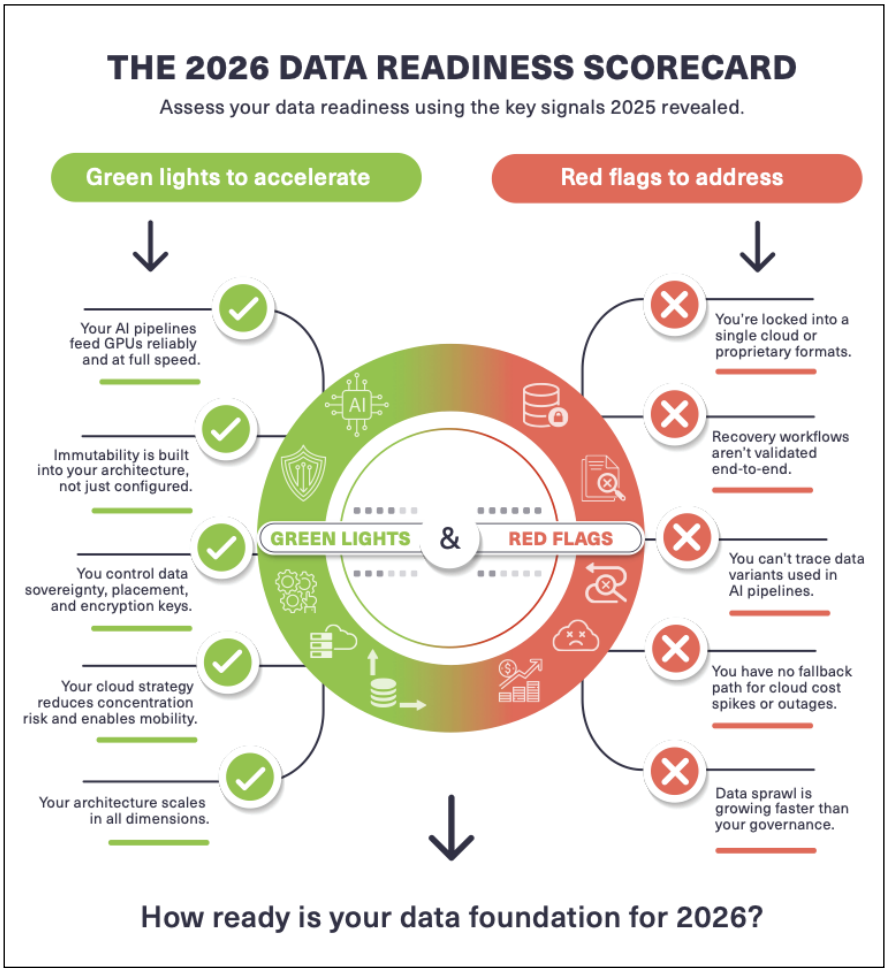

Ignore at your own risk: The hidden price tag of bad code

Bad code carries a cost that touches budgets, security, operations, employee wellbeing, and long-term competitiveness. However, these costs are avoidable through pairing the creativity and speed of AI tools with the expertise and critical thinking of trained developers.

BY BEN MATTHEWS, SENIOR DIRECTOR OF ENGINEERING AT STACK OVERFLOW

AI TOOLS HAVE democratised the power of software, bestowing the novice with the ability to generate functioning code in moments. The emergence of vibe coding captures the excitement felt by non-technical users who suddenly find themselves capable of building applications - a milestone all developers can remember - now through conversational prompts rather than structured learning. Yet, beneath this sense of empowerment lies a growing risk for businesses. While AI can generate code at speed, it cannot guarantee its quality for the application scenarios they are targeting. For many organisations, this ease of creation is ushering in a new layer of technical fragility - one that carries tangible operational and financial costs.

When quick fixes cover deeper flaws

The rise of AI-generated coding has led to a rise in software that appears functional but can often fail under pressure. Code produced through automated tools often generalises architecture, ignores bespoke needs, and can lack context for security concerns that your app may introduce; things that senior developers know and value through time tested experience. This pattern reflects a growing problem across industries: code that looks correct on the surface can mask a foundation that lacks the base that traditional development and QA tools provide.

Organisations that adopt AI-generated development without proper stewardship from engineers who understand what the code is doing are discovering that these shortcuts inevitably lead to bottlenecks and time sinks in other areas. What began as an efficient solution transforms into a long-term liability that nets a need for more scrutiny, developer time and overhead to address gaps.

The rising financial burden of poor software

Bad code is not just a technical concern - it is one of the most persistent and underestimated financial drains in modern organisations. Often called technical debt, developers have long warned that the cost of fixing code defects grows exponentially the later they are addressed or, even worse, discovered. As a result, much of the most valuable engineering knowledge today is summarised rather than explored and understood.

Even when bad code does not lead to system failure, its inefficiencies can quietly escalate financially as LLMs build things without quality or cost in mind behind the scenes. One of the ironic lessons that junior engineers learn early on is that fatal errors are often better than non-fatal ones simply because they are easier to identify; teams spend much more time on trying to track down errors that fly under the radar than the obvious ones. Organisations absorb higher operating expenses without gaining the competitive advantage that well-engineered software is designed to deliver if they can't find problems, are building on an unstable base and

are constantly trying to realign their applications.

Security risks hidden in plain sight

The most dangerous consequence of low-quality code is the security lapses it can create. AI-generated code often pulls patterns from training data without consideration of security best practices. Input validation may be weak or non-existent. Authentication steps may be oversimplified and outdated dependencies may be silently included. These vulnerabilities become opportunities for bad actors - and because the code appears functional, many organizations fail to detect the risks until it is too late, such as this instance where vibe coding produced an account verification email that contained the actual passcode in the email. While it looks correct at first glance, the AI tool immediately compromised security.

Security teams then face the challenge of reactively fortifying systems that were never built with protection in mind and constantly retrace the AI's steps in a constant defensive posture. This approach is expensive, time-consuming, and insufficient in a threat landscape defined by rapid exploitation and sophisticated attack chains. Using AI to learn and react to threats makes sense, but using it as a base from which you build can be an expensive path.

The human cost: Burnout and bottlenecks

Developers feel the impact of bad code more than anyone. In organizations where rapid output is prioritised over reliable expertise, engineers become trapped in a reactive cycle. Instead of building new features or exploring innovation, they spend their time stabilizing systems that were created hastily or without technical understanding. If you think building at a steady pace is expensive, wait until you see how long it takes when you build it too fast.

This environment is a direct contributor to burnout. Repeatedly solving preventable problems erodes morale, and teams will lose the space to think strategically, improve architectures, or propose meaningful improvements. As technical debt grows, engineering velocity slows. Eventually, the organization finds itself restrained not by a lack of ideas, but by the weight of past decisions.

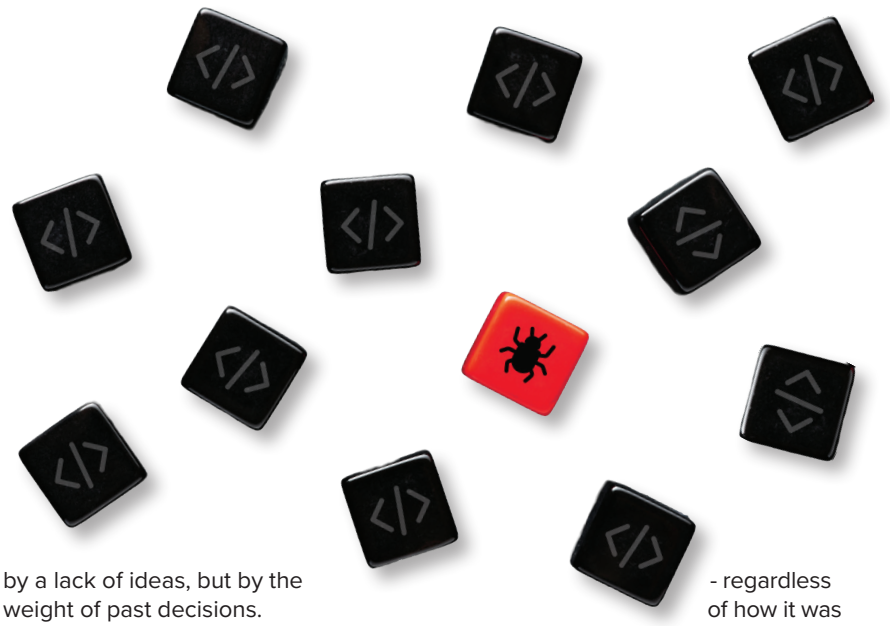

Why skilled developers still matter

Despite fears that AI will replace developers, the value of experienced engineering has increased in an era where "anyone" can write code. Real developers do more than produce syntax; they anticipate edge cases, plan for scalability, embed security principles, and build systems designed to evolve. Not only do they know how to write code, but they can also recognize which patterns are safe, scalable, and battle-tested. One study suggests that developers who know how to instruct AI receive higher accuracy scores, with scores jumping as much as 17.7% to 78.7% in some cases. The value of this experience isn't just hand writing code, but in using and understanding how a new tool functions.

AI can accelerate tasks, support prototyping, and inspire creativity. However it cannot replicate intuition, judgement, or the ability to foresee downstream consequences that AI needs to function correctly. As more people gain the ability to generate code, the role of curated, experience-driven knowledge becomes even more important. Organizations that treat AI as a shortcut rather than a tool risk undermining the very foundation of their platforms.

Building a future where speed doesn't compromise quality

The path forward for businesses is not to reject AI-assisted coding, but to implement it with discipline. Governance frameworks must ensure that all code - regardless of how it was generated - is reviewed, tested, and validated by people with the expertise to judge its reliability. Engineering standards should remain uncompromised and the culture should reward careful design as much as rapid experimentation. The goal is not fewer questions, but fewer avoidable mistakes leading to bad code development.

Organizations that strike this balance will gain the benefits of AI, which can be plentiful, without falling victim to its pitfalls. They will develop faster while maintaining the integrity of their systems. They will innovate without collecting destructive technical debt. Most importantly, they will empower their engineering teams instead of overwhelming them. In doing so, they transform shared technical knowledge into a strategic asset rather than an afterthought.

Quality remains the competitive advantage

Bad code carries a cost that touches budgets, security, operations, employee wellbeing, and long-term competitiveness. However, these costs are avoidable through pairing the creativity and speed of AI tools with the expertise and critical thinking of trained developers. Reliable, community-validated knowledge plays a crucial role in ensuring that speed does not come at the expense of quality. The companies that succeed in the AI era will be those that understand that while anyone can generate code, building *good* software still demands skill, foresight and a commitment to excellence.

Why LLMs are plateauing – and what that means for software security

There's no doubt the AI-generated code landscape evolved at an unprecedented rate over the last year. The rise of vibe coding, where developers use large language models (LLMs) to generate functional code, has fundamentally changed how software is built.

BY JOHN SMITH, CTO EMEA AT VERACODE

AS AI BECOMES embedded into everything from apps to full-scale company operations, it's clear significant effort has gone into training LLMs for correctness, with newer and larger models becoming increasingly effective at generating code with the expected functionality. Less attention, however, has been paid to whether the produced code is secure. The result? Mountains of production code that works in practice but is quietly embedding and spreading significant security vulnerabilities.

At the same time, LLMs are enabling attackers to identify and exploit these flaws faster than ever. With defence capabilities lagging behind, the gap between attackers and defenders is widening at a critical time, just as enterprises are increasingly reliant on AI.

Security is flatlining across most AI models

Despite recent rapid surface-level progress in AI models and their ability to generate functional code, there is growing evidence that security has failed to keep pace. Recent research shows most generative AI (GenAI) tools are producing glaring security flaws, including popular models such as Anthropic's Claude, Google's Gemini, and xAI's Grok. Across all models, languages, CWEs (common weakness enumeration) and tasks, only around 55% of generation tasks produce secure code, meaning LLMs are introducing a detectable OWASP Top 10 vulnerability nearly half the time.

Surprisingly, this heightened vulnerability risk was agnostic across the different model types, with no significant difference between the smaller and larger models. Whilst the ability to generate syntactically correct code has improved dramatically, security remains stubbornly stagnant. Simply scaling models or updating training data is insufficient to meaningfully improve security outcomes.

The notable exception is OpenAI's reasoning GPT-5 models, which take extra steps to think through problems before producing code. These models achieved substantially higher security pass rates of 70% and above, compared to 50-60% for previous generations. However in contrast, GPT-5-chat, a non-reasoning variant, lagged at 52%, suggesting that reasoning alignment, not model scale, drives these gains. It's possible OpenAI's tuning examples here include high-quality secure code or explicitly teach models to reason about security trade-offs to achieve this higher rate.

Language-specific trends have also emerged. Many of the AI models perform much worse on Java code generation tasks than any other coding languages, with security pass rates at less than 30%, while Python, C# and JavaScript generally fall between 38% and 45%. At the same time, newer models, especially reasoning-tuned ones, are performing better at

generating secure C# and Java code, likely reflecting AI labs' focus on major enterprise languages.

Why is LLM security stagnating?

The root of the problem lies in the nature of the training data, made up of public code samples scraped from the internet. As a result, the data contains both secure and insecure examples, including deliberately vulnerable projects like WebGoat – an insecure Java application used for security training. The models then treat all these examples as legitimate ways to satisfy a coding request, learning patterns that don't reliably distinguish safe from unsafe implementation.

With most LLMs training on this publicly available data, there are similar patterns in how they produce security risks. As the data remains largely unchanged over time, and is increasingly supplemented with synthetic and AI-generated code, model security performance has remained stagnant across generations of models.

This also helps explain why Java is particularly problematic. Java has a long history as a server-side implementation language, and predates widespread recognition of vulnerabilities like SQL injection. Its training data must therefore contain many more security vulnerabilities than other languages like C# or Python, leading models to perform significantly worse on Java-specific tasks.

The security blind spot in vibe coding

These findings raise huge concerns for AI-assisted development and the growing popularity of vibe coding.

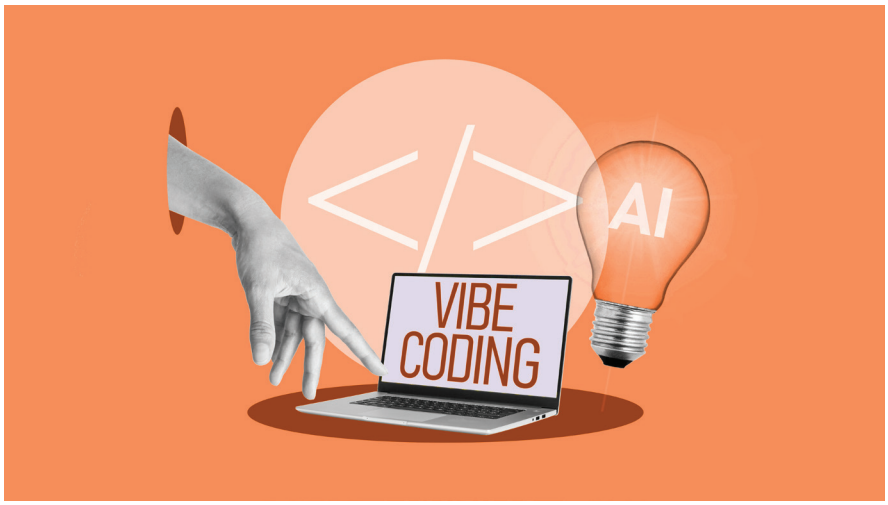

While these practices accelerate productivity, developers rarely specify security constraints when prompting LLMs, which would dramatically improve the security of generated code.

For example, a developer might prompt a model to generate a database query without specifying whether it should construct using a prepared statement (safe) or string concatenation (unsafe). This effectively leaves those decisions down to the LLMs which, as the findings show, choose incorrectly nearly half the time. Alarmingly, this issue shows little sign of improving.

And the risks are already surfacing in practice. A recent incident with an AI coding tool on the Replit platform caused the deletion of an entire live production database during a code freeze – a clear warning of what can go wrong when AI-generated code is trusted without sufficient guardrails.

The implications for developers and organisations

Given these persistent shortfalls, relying on model improvements alone is not a viable security strategy. While newer reasoning models offer a clear advantage, security performance remains highly variable and even the best-performing models introduce vulnerabilities in nearly a third of cases.

AI coding assistants are powerful tools, but they cannot replace skilled developers or comprehensive security programmes. A layered approach to risk management is essential: maintaining continuous scanning and validation using static analysis (SAST) and software composition analysis (SCA), regardless of code origin, and proactive

Java has a long history as a server-side implementation language, and predates widespread recognition of vulnerabilities like SQL injection. Its training data must therefore contain many more security vulnerabilities than other languages like C# or Python, leading models to perform significantly worse on Java-specific tasks.

blocking of malicious dependencies are crucial to preventing vulnerabilities from reaching production pipelines.

AI-powered remediation tools also assist developers by providing real-time guidance and automated fixes, yet responsibility for secure implementation ultimately remains human.

The hidden cost for AI generated code

AI coding assistants and agentic workflows represent the future of software development and will continue to evolve at a rapid pace. But while LLMs have become adept at generating functionally correct code, they continue to produce security vulnerabilities at a troublingly high rate – an issue that won't be easy to fix.

The challenge for every organisation is ensuring security evolves alongside these new capabilities. Addressing this requires security-specific training, reasoning alignment, and a recognition that security cannot be an afterthought if we want to prevent the accumulation of masses of security debt.

Until AI labs prioritise security in training and alignment processes, developers and security teams must treat AI-generated code as an inherently untrusted input – a principle that must be considered in day-to-day vibe coding.

Why software is the essential building block behind quantum computing's huge potential

Quantum computing is no longer a distant concept. It's rapidly becoming a strategic priority for multiple industries seeking breakthroughs in materials science, energy efficiency, molecular modelling, and drug discovery. However, even the most advanced quantum hardware is ineffective without robust software, and today, quantum software remains far from mature.

BY ERIC PAULSEN, FIELD CTO, EMEA, CODER

THE SHIFT WE'RE seeing is not in the physics behind quantum, but rather the urgency to find workable applications. As fault-tolerant quantum systems move closer to feasibility, organisations are asking what quantum development looks like in practice. And the reality is that most of the procedures and approaches that developers learnt from classical computing just don't apply.

Trial and error are not in the quantum vocabulary

Consider iteration. In traditional software development, iteration is fast and inexpensive. Developers write code, test it, debug it, and repeat the process until they get it right. Quantum development is fundamentally different, and does not allow for this trial–and-error methodology. To start, developers cannot run a quantum program on a laptop, and simulating even modest qubit counts (the quantum bits that process information) is computationally prohibitive. 30 qubits is about as much as a supercomputer can handle and by 50 qubits the world doesn't have enough classical computing power to make the simulation worthwhile. This constraint inevitably forces a rethink of development workflows.

Hardware challenges compound the issue. Early quantum machines were delicate and noisy, often failing after a few hundred operations, making calculations entirely meaningless. While today's fault tolerance improvements offer stability through error correction, they also introduce complexity. As quantum systems scale, control mechanisms become increasingly intricate—and significantly harder to secure.

The solution lies in abstraction. It is too much to expect developers to manually control qubits, noise, correction logic, and security. Just as classical computing evolved beyond direct memory management and tweaking hardware to fit the task in hand and became automated and efficient, quantum software must follow suit. Emerging software toolchains are already reducing complexity, guiding developers toward proven patterns and automating low-level decisions. Libraries now handle qubit allocation, circuit design, and resource tracking—critical software-enabled steps toward making quantum development accessible beyond those with a physics degree.

The need for a secure development environment

Security, however, remains a major concern. Quantum research intersects with national strategic interests, making algorithms and architectures high-value targets. The risk of encrypted data being harvested today and decrypted at a later point is real, reshaping what secure development environments require. For many organisations, loosely managed or shared platforms are no longer viable, and it is only systems with end-to-end control that will pass the acceptance test. Data sovereignty, strict access policies, and integration with enterprise security frameworks are non-negotiable. Adopting quantum computing cannot come at the price of reducing security protocols.

PsiQuantum exemplifies this approach. Rather than adapting existing platforms, the company — which is US-based and develops million-qubit, fault-tolerant quantum computers — has built secure, preconfigured environments tailored to quantum's unique constraints. Researchers and developers operate within isolated toolchains on infrastructure that the organisation can directly control, which ensures that sensitive work remains protected while operational friction is kept to a minimum. This model is gaining traction as businesses prepare for quantum systems capable of solving real-world challenges within the next ten years.

Reaping the collaboration effect

Despite the security imperatives, open source remains essential to the evolution of quantum computing. Many of today's quantum tools, such as compilers, simulators, and resource estimators, exist because of open and collaborative development. Organisations hoping to harness the advantages of quantum may find that balancing openness with security is challenging, but this shouldn't be a barrier. Across the world, businesses are finding ways to combine open tooling with private workflows rather than keeping the two mutually exclusive.

The fact is that the future success of quantum computing does not lie in hardware alone. The environments and software development tools that allow teams to experiment and collaborate securely are equally critical. Organisations investing in these capabilities today will be positioned to lead when quantum machines are fully ready to deliver on their promise.

The next era of computing will be defined not only by new physics, but also by how effectively we enable development in fit-for-purpose environments so systems can be both usable and secure.

Save THE Date

MSP CHANNEL AWARDS 2026

CELEBRATING 17 YEARS OF SUCCESS

SAVE THE DATE: 26 NOVEMBER 2026

A new chapter is coming for the MSP Channel Awards.

Now in its 17th year, the programme continues to celebrate excellence across the channel and the incredible work being done across the industry. As the sector continues to evolve, the awards evolve with it, reflecting the innovation, growth and impact shaping the channel today.

Managed Service Providers are now at the heart of delivering innovative, cutting-edge IT solutions, and this year's awards will once again shine a spotlight on that impact. Alongside the core categories you know and recognise, there will also be new and exciting categories introduced this year to better reflect the breadth and diversity of the ecosystem.

This is your opportunity to be recognised for the work you are doing and the impact you are making.

You can submit as many products or projects as you like, giving you the opportunity to showcase your innovation and highlight the work you're most proud of.

Nominations will be opening soon, so keep an eye out and make sure you do not miss your chance to get involved and put your achievements forward.

mspchannelawards.com
T: +44(0)2476 718 970

Angel BUSINESS COMMUNICATIONS

SDC AWARDS

Unlocking AI's true potential: why high-quality first-party data and orchestration drive real business value

By investing in DEX platforms and solid data management tools, organisations can dodge the usual AI failures and actually tap into what it can do.

BY DAN SALINAS, COO, LAKESIDE SOFTWARE

ACROSS LARGE ENTERPRISES, AI initiatives rarely collapse in obvious ways. More often, they stall quietly. Pilots look promising, models perform as expected, but the day-to-day reality doesn't change much for employees or IT teams.

When that happens, the issue is usually not the intelligence itself, but the surrounding system. Fragmented data and disconnected workflows make it difficult to turn AI insight into action at scale. At the centre of this challenge is orchestration: coordinating data, systems, and actions so AI can consistently deliver meaningful outcomes, not just insights.

In this article, we look at why that gap persists and what organisations are missing. We break down the three elements that consistently distinguish AI programmes that deliver real value from those that don't: high-quality first-party data, system flow, and reliable output. We also explore how orchestration connects these pieces and, for CIOs and COOs, why this represents an operating model decision that directly affects productivity, cost, and risk.

The foundation: high-quality first-party data to ground AI in reality

Everyone knows that AI does not compensate for weak data. Poor-quality or generic inputs simply produce equally weak outputs, unclear responses, and outright errors. First-party data, gathered directly from an organisation's endpoints and systems, addresses this problem by providing a real-time, proprietary foundation that prevents AI from relying on unhelpful, broad generalisations.

Digital Employee Experience (DEX) monitoring platforms, such as SysTrack AI, exemplify this approach. They gather information from multiple endpoints such as devices, applications, networks, and user behaviours, even when they are offline. This high-frequency data is not only extensive; it's also structured and contextual, forming the backbone of effective AI implementation.

According to a Gartner survey, 85% of AI projects fail due to poor-quality data. Beyond that, the same Gartner report found that 63% of organisations either lack or are unsure whether they have the right data management processes for AI projects.

Without a solid foundation, AI systems risk "hallucinations" based on incomplete or biased training data. However, first-party data changes this

dynamic, with AI agents that work with enterprise-specific information to provide contextual insights. For instance, in a global healthcare insurance provider we worked with, this data helped proactively identify potential issues, reducing 30,000 incidents across their organisation and cutting downtime that could affect patient care.

The engine: flow for seamless automation and process efficiency

Data at rest provides little value, so that's where flow comes in. Flow transforms static information into dynamic workflows, enabling AI to handle tasks automatically and removing interface friction that slows business operations. More importantly, flow enables organisations to redesign how work moves across teams and systems, not just speed up individual IT tasks.

In practice, this means integrating AI with existing tools such as IT service management (ITSM) systems, workflow engines, and APIs to enable efficient actions. AI-augmented IT operations and diagnostics technologies, such as SysTrack AI, demonstrate this in operation. The platform collects telemetry, identifies patterns, diagnoses root causes with high confidence, and acts based on persona-aware instructions. Whether delivering a quick fix for end users or generating a detailed ticket for IT teams, the system adapts its response to context.

The outcome is significant efficiency gains. Organisations implementing these flows report reductions of up to 25% in help desk ticket volumes and 75% in mean time to resolution (MTTR). Those gains come not from replacing human expertise but from automating repetitive diagnostics and allowing teams to focus more on complex challenges.

Reliable output for content, code, and action

The final element is output. AI delivers value when it produces accurate and understandable results that users can trust and action. When output is built on high-quality first-party data and supported by smooth flows, it becomes repeatable and verifiable, reducing errors and escalations.

Explainable reasoning is central to this reliability. Continuous telemetry provides transparency at each step. For example, diagnosing a performance slowdown, providing step-by-step solutions, and confirming fixes, all while directing information to the appropriate persona, whether that's an L3 agent for complex IT issues or a self-service portal for employees.

This reliability is crucial in high-stakes situations. When IT solutions are safe and understandable, organisations can automate across global teams, using specific insights to boost productivity without introducing new risks.

Orchestration as the intelligent traffic controller

Individually, each of these elements, data, flow, and output, has its own limits. Data that doesn't move creates absolutely zero value. Processes that move but produce unreliable results generate confusion. Outputs based on questionable data erode trust. This is where orchestration connects the elements, functioning like a smart traffic controller that interprets user needs, routes tasks to the right systems, retrieves relevant data, triggers appropriate workflows, and ensures proper delivery.

Consider a managed service provider managing over a million seats. Improving workflows with endpoint context requires orchestration to blend data from the edge with human know-how. The outcome is better service delivery and a better employee experience. Problems are resolved faster, costs go down, and teams become more productive.

Companies are racing to adopt AI, but realising true return on investment depends on prioritising high-quality data and orchestrating it effectively. That's what unlocks actual value, boosts productivity, maintains security, and drives innovation. CIOs and COOs need to own this, making sure AI isn't merely deployed but properly set up to succeed.

By investing in DEX platforms and solid data management tools, organisations can dodge the usual AI failures and actually tap into what it can do. The future isn't about having more AI. It's about smarter orchestration driven by quality data that delivers measurable results. That's what separates experiments from transformation.

Does power consumption really offset the cost gap between Flash and HDD? Absolutely not

There is a persistent claim in the storage world that flash is worth its higher price as it consumes less power. The idea sounds appealing, especially in an era of rising energy costs and sustainability targets. Many vendors highlight lower wattage as a way to justify premium flash systems and to soften the impact of rising drive prices. But the truth is simple: power savings do not come close to offsetting the massive cost gap between flash and HDD. Not now, and not in any realistic future scenario.

BY GAL NAOR, CEO, StorONE

The price gap is much wider than any energy savings

FLASH PRICING has climbed sharply due to component shortages, manufacturing constraints, and overwhelming demand from AI and large-scale cloud environments. HDD prices, meanwhile, remain stable and far more predictable. The difference between the two has widened to the point that energy efficiency is irrelevant compared to the upfront cost of the media itself. Even in data centers with high electricity rates, the total cost of powering a drive over its entire lifespan is a fraction of the cost of the drive. Energy consumption simply does not move the overall budget enough to change the purchasing decision. Electricity is the rounding error. Storage hardware is the actual expense.

Flash does not pay for itself through power savings

To understand why, it helps to think about another common household comparison. Many people replace air conditioners hoping to cover the new unit costs by the reduction in their electricity bill, but an AC unit is a heat pump that draws a predictable amount of energy regardless of its age. The electricity savings from a newer AC are marginal unless the old system was failing. In reality, people replace ACs for comfort, reliability, or functionality, not because they expect dramatic utility savings. The same logic applies to flash. Lower wattage does not magically cover the significant premium for the drives. Flash is purchased for performance, not because it offers a meaningful cost-return model on electricity.

Solar panels create ROI. Storage devices do not

Some people compare flash to energy-efficient technologies like solar panels. But solar panels eventually produce energy at zero marginal cost, allowing the owner to recoup the initial investment. Storage media does not behave that way. When flash wears out, it must be replaced at full cost. Whatever tiny amount of energy you saved over its life disappears the moment you purchase its replacement.

There is no cumulative benefit. There is no long-term return. The economics reset to zero with every new drive.

HDDs are far more efficient than their reputation suggests

HDDs are often perceived as power-hungry, but modern high-capacity drives are surprisingly energy-efficient. Their wattage per terabyte has improved dramatically, and their overall operating cost remains low. When evaluated on a per-dollar or per-petabyte basis, HDDs remain unmatched in cost efficiency.

The narrative that HDDs are outdated and inefficient has been repeated so often that it is rarely questioned. Yet the numbers consistently show that HDDs offer strong performance per watt and a far lower cost per usable terabyte. In contrast, flash reaches end-of-life much sooner and must be replaced far more frequently, increasing total lifetime cost.

The only scenario where power truly matters

There is one exception. If a facility genuinely has limited power or cooling capacity, the choice of media becomes an engineering constraint. This may force architectural changes that favor lower-wattage technology. But such design will not be optimal from the TCO perspective. It is simply a logistical reality of an undersized or fully utilized data center. Most organizations are nowhere near that limit. They are making financial decisions based on budget, not on whether they have reached their power envelope.

Power consumption does not and will not bridge the cost gap between flash and HDD

The difference in upfront price is simply too large for energy savings to matter. Flash will reach its endurance limits long before any savings come close to recovering the premium. HDDs remain efficient, reliable, and cost effective for cold and warm data, and their energy footprint is far smaller than industry myths suggest.

As the storage industry faces rising flash prices, unpredictable component availability, and accelerating data growth, organizations must focus on true economic efficiency, which starts with recognizing that power savings are not the justification for preferring flash over HDD. Storage decisions should be based on workload behavior, data temperature, and long-term scalability, not on a myth that energy savings can offset expensive hardware.

EXAGRID®

Tiered Backup Storage

The future is here.
Tiered Backup Storage

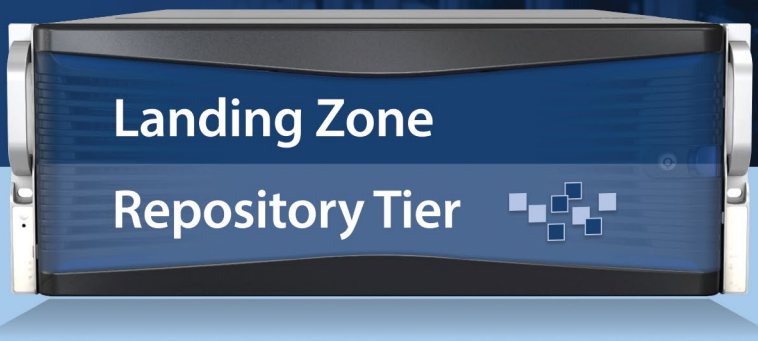

FASTEST BACKUPS

FASTEST RESTORES

SCALABILITY FOR FIXED-LENGTH BACKUP WINDOW

COMPREHENSIVE SECURITY WITH RANSOMWARE RECOVERY

LOW COST UP FRONT AND OVER TIME

MSP CHANNEL AWARDS 2025 WINNER

- BACKUP & DR INNOVATION OF THE YEAR
- STORAGE HARDWARE INNOVATION OF THE YEAR
- STORAGE VENDOR OF THE YEAR

Thank you so much to all who voted, and congratulations to our fellow MSP Channel Awards 2025 winners!

Visit our website to learn more about ExaGrid's award-winning Tiered Backup Storage.

LEARN MORE >