ENTERPRISE AI

ENTERPRISE AI

Distributed in
THE TIMES

Contributors

Sam Birchall
Lead writer at Raconteur, focusing on the inner workings of the finance function and the trends shaping the future of financial leadership.

Tom Dennis
Editor and journalist covering the forefront of tech and innovation for almost two decades.

Joy Persaud
Freelance journalist with over a decade of experience covering business, health, education and lifestyle topics.

Raconteur

Editor
Tom Dennis

Commercial content editor
Jessica Lynn

Commercial production managers
Alex Datcu
Audrey Davey
Ellen Newsome

Design and illustration
Kellie Jerrard
James Lampard
Celina Lucey
Samuele Motta

Design director
Tim Whitlock

HUMAN RESOURCES

The AI job carnage is taking a mental toll on workers

Sleepless nights, lost work identity and resistance to change: the fear of being replaced by AI is taking effect in the workplace

Sam Birchall

Replacement is now weighing heavily on employees. Research by KPMG shows that 52% of workers believe AI will harm their job security, fuelling anxiety — a narrative that is being further pushed by senior executives. Anthropic's CEO, Dario Amodei, warned that AI could wipe out half of all entry-level white-collar jobs and Microsoft's AI CEO Mustafa Suleyman said that AI could automate "most, if not all" white collar tasks within 18 months.

Experts are now warning of the potentially devastating psychological toll that AI automation, or the mere threat of it, can have on the workforce. In a recent academic paper, two researchers argue the phenomenon is significant enough to merit its own clinical label: AI replacement dysfunction, or AIRD. According to the authors, the constant fear of job loss could be driving symptoms ranging from anxiety, insomnia, paranoia, and loss of identity, even among otherwise healthy individuals.

So far, much of the debate around AI and mental health has focused on the personal risks of using the technology itself, including reports of chatbots fuelling delusions or encouraging harmful behaviour. But the broader emotional impact of simply living and working under the shadow of AI-driven displacement has been largely overlooked. As headlines about AI-driven layoffs mount, AIRD may demand far greater attention from employers.

Even without layoffs, rapid technological change can spark anxiety among staff. Dr. Brittany Straton, senior lecturer in cyberpsychology at Arden University, calls it a "chronic stressor" that erodes wellbeing, psychological safety, and motivation. "Our jobs give many of us a sense of purpose and identity," she explains. "When tasks are automated, workers can experience a loss of professional identity or purpose."

The irony is that AI doesn't need to replace people to cause strain. The mere perception of risk can influence behaviour, Straton explains. Workers may withdraw, resist new technology, or even hide knowledge to protect their perceived value. "These are not signs of unwillingness to modernise – they're stress responses rooted in identity protection," she says.

Performance pressures compound the strain, according to Tracey Paxton, clinical director at Perkbox, an employee benefits group, and chair of The Royal College of Psychiatrists APPTS Advisory Board. "Some workers feel they must compete with machine efficiency or constantly upskill to survive, creating a sense of never being secure or good enough. This is a known driver of stress and presenteeism," she says.

Some studies suggest that reducing low-control, high-repetition tasks can improve both wellbeing and engagement. "While some employees feel genuine relief, others experience a subtler, corrosive stress," Paxton explains. "Uncertainty is one of the strongest drivers of workplace stress, and AI can intensify it — raising questions about role stability, skills relevance and long-term employability."

In Paxton's view, when changes are introduced quickly or with limited staff involvement, trust in leadership can decline, with employees feeling replaceable or overly monitored. "This combination — reduced control, weakened role identity, and lower trust — is exactly what organisational change research predicts when technological transformation is experienced as imposed rather than collaborative."

The good news is employers are increasingly aware of these psychological effects. "I'm seeing more transparent communication about how AI will be used," Paxton says. "Where organisations involve staff early, evaluate progress and emphasise augmentation rather than replacement, employee fear reduces significantly."

Dr Aaron Taylor, head of human resource management at Arden University, is seeing more emphasis on training and reskilling as core elements of AI strategy. "Employers are investing in structured upskilling programmes, coaching and digital literacy support," he says. "These initiatives build capability while reinforcing a sense of future employability, which is vital for morale. When workers can see a pathway through the change, their confidence and engagement increase."

Other firms are pairing technological changes with wellbeing checks and manager training on how to talk about AI. "Workers are looking for the truth about what AI will change, reassurance about where humans remain essential, and support to grow into new forms of work," Taylor says.

By failing to address the psychological impact of AI on their workforce, business leaders may be exposing themselves to legal risk. "Workplace AI anxiety can become a liability once an employer is, or reasonably should be, aware of a genuine risk of harm and fails to take steps to mitigate it," says Hannah Mahon, a partner in Eversheds Sutherland's Employment, Labor and Pensions group.

While litigation directly linked to AI adoption is still emerging, existing legal duties around psychiatric harm, workplace stress and job security already apply in many cases, Mahon explains.

HR and legal teams therefore play a critical role. Documenting employee concerns, responding promptly, and maintaining open communication channels not only builds trust but also demonstrates that employers have taken reasonable steps to safeguard staff.

"Sharing information early, explaining AI objectives in clear, accessible language, and involving employees in decision-making helps people feel included rather than sidelined – reducing stress and the likelihood of grievances," Mahon says.

Although UK law does not yet explicitly address AI-related psychological harm, employers remain bound by established duty-of-care obligations to manage foreseeable risks arising from workplace technologies. According to Mahon, existing frameworks covering fairness, discrimination and employee wellbeing are broad enough to encompass AI-related challenges, though more targeted guidance may emerge as the technology evolves.

In practice, minimising both harm and liability comes down to proactive engagement, transparent communication and positioning AI as a tool for augmentation rather than replacement – approaches that protect employees while helping organisations navigate the transition responsibly.

JOBS EXPOSED TO AI

AI Job Risk Index, 2026

The likelihood that different professions will be automated by artificial intelligence, average risk score (out of 100)

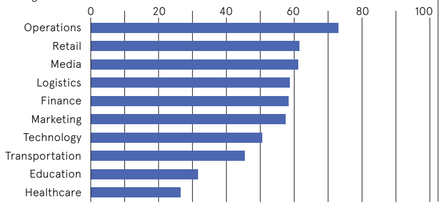

Commercial feature

Zero trust for humans – but implicit trust for machines?

The rise of AI agents is testing zero trust security, exposing risks that organisations can neither fully see nor reliably control. In response, enterprises must strengthen visibility and governance

New research has revealed the extent to which the zero trust framework, developed to reduce risk across enterprises, is under pressure as AI adoption outpaces security governance. In 65% of organisations, zero trust controls cannot secure non-human identities (NHIs), including new agentic AI systems.

AI agents offer clear advantages, from generating content and retrieving information, to triggering downstream actions. However, they often run unsupervised and without guardrails, increasing the risk of data leaks, credential compromise and wider operational disruption. The fact that they are operating with fewer checks than their human colleagues should ring alarm bells.

According to Netskope's *AI Risk and Readiness Report 2026*, which surveyed 1,253 cybersecurity professionals, 56% of enterprises acknowledge exposure to agentic AI risk. This is largely because AI tools operate autonomously in shadow mode, with organisations often only discovering what an agent has done after the action is complete.

Shadow AI risk grows

The scale of adoption is already significant: some 24% said agents were in limited production within their organisation, 9% had ungoverned agents operating at scale handling core business logic, and 23% suspected there to be shadow agentic AI deployments in operation, unknown to internal IT. In fact, 32% admitted that they have no visibility into agent actions at all.

"Organisations need a better understanding of the underlying technology and greater visibility into what they are giving up when they use agentic AI," says Netskope's CISO, James Robinson. "Too many enterprises are relying on legacy security models to secure this new technology."

AI agents often have broad access across enterprise systems and almost none of it can be meaningfully intercepted. They can also be prompted to perform unintended actions – for example, reacting to malicious prompt injections embedded in open-source software or external data sources. In such cases, an agent may follow instructions that appear legitimate but have been manipulated to trigger harmful behaviour.

Once an agent initiates a harmful action, only 9% of organisations can intervene before it completes, according to Netskope. Of that figure, 24% can block some actions but not all, 35% only identify them in logs after completion, and 32% have no visibility at all.

This lack of control highlights a fundamental shift. So what is the future of the zero trust framework in a world where 91% of organisations cannot stop an agent before it acts?

Rethinking zero trust

Zero trust was built around a human user – with a device, location, behavioural pattern and risk score. By contrast, an AI agent has a credential, a defined scope and a task. It does not behave like a person, nor does it follow predictable patterns of intent. The principles of zero trust must therefore evolve to account for non-human identities and machine-to-machine interactions, where actions are faster, more opaque and harder to audit.

The Netskope study shows these risks are already materialising. Some 37% of respondents report AI agent-related operational issues in the past year, with 8% resulting in outages or data corruption. These are not theoretical concerns but active failures already affecting production environments.

Simply banning the use of agentic AI is not the answer. It risks driving adoption underground, making it harder to govern and contain when problems arise. As with previous waves of technology entering the workplace, shadow usage is often more dangerous than controlled deployment.

Some 62% of organisations already apply zero trust principles to AI security in some form, largely based on CISA's five pillars: Identity, Devices, Networks, Applications and Workloads and Data. While these remain relevant, the Identity pillar is under particular strain as legacy architectures are applied to non-human identities that do not fit traditional models.

The challenge becomes more complex when individuals deploy multiple agents. Should each agent have its own identity, permissions and audit trail, or inherit that of the human directing it – for example, when answering emails, accessing documents or responding to queries? The wrong approach risks either over-privileging agents or losing accountability entirely.

Defining accountability

Ultimately, says Robinson, enterprises must define how much risk they are willing to tolerate.

This requires employees to understand their responsibilities, and organisations to clearly define acceptable use. It also raises a critical question: how much visibility should an organisation have into its agents in case one goes rogue and causes operational damage?

Robinson says enterprises must strengthen how they monitor AI agents and, as tools mature, develop the ability to intercept actions at the request layer. This means moving beyond traditional perimeter-based controls and towards real-time inspection of what an agent is being asked to do and how it responds. Zero trust access controls must extend to securing non-human identities with the same rigour applied to human users.

"You need to define what anomalous looks like for agent behaviour in your environment, build detection rules for those patterns and require human-in-the-loop approval for high-risk actions such as account creation, permission changes and external data transfers," he says.

He also urges security professionals to better understand the evolving threat landscape and to treat incidents as opportunities to strengthen systems. Rather than viewing failures as isolated events, organisations should use them to refine controls, improve visibility and reduce the likelihood of recurrence. As AI agents become embedded across productivity platforms, organisations must not only fix issues but ensure systems improve as a result.

It is also vital to have vigorous incident response protocols in place so that forensic investigations can be carried out effectively.

"Organisations should adopt new security architectures and see security not as a barrier but as a way to enable safe innovation within defined guardrails," says Robinson. He also recommends running internal "AI prompt-athons" where teams can safely explore how agents behave, test edge cases and better understand potential risks in real-world scenarios.

There is no doubt enterprises are increasing investment in both AI and the systems to secure it – 90% have increased their AI security spending in the past 12 months. But confidence in the technology and understanding of its risks remain relatively low. This gap between adoption and preparedness is where many of the current vulnerabilities are emerging.

Robinson sees a significant need for organisations to bring together technology, people and processes to close the visibility gap and implement effective agentic AI frameworks, systems and guardrails. This includes extending activity-level monitoring, distinguishing between personal and corporate AI accounts and ensuring that machine-to-machine interactions are subject to the same scrutiny as human activity. Every AI deployment carries risk, whether organisations recognise it or not.

56% of organisations acknowledge exposure to agentic AI risk

91% of organisations cannot stop an agent before it acts

Netskope, 2026

> The principles of zero trust must evolve to account for non-human identities

BUSINESSES ARE BETTING BIG ON AI

As AI adoption accelerates, firms are seeing measurable gains in productivity, decision-making and customer engagement. But the struggle to turn advances into tangible revenue remains a major challenge. As the technology is further integrated into business operations, questions over strategic alignment, and the true impact on long-term growth are coming to the forefront, forcing executives to rethink how AI shapes their competitive advantage

TOP BENEFITS ACHIEVED FROM ENTERPRISE AI ADOPTION

Based on weighted responses of global leaders reporting tangible gains

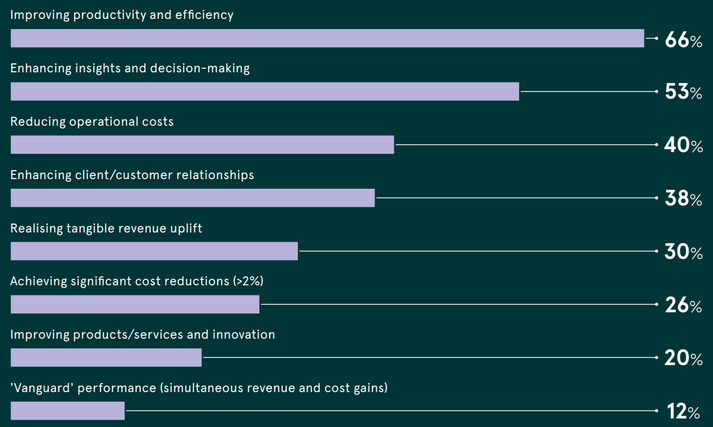

Deloitte, 2026

GLOBAL INVESTMENT HOTSPOTS FOR ENTERPRISE EXPANSION IN 2026

Percentage of CEOs identifying top three countries for highest investment proportion

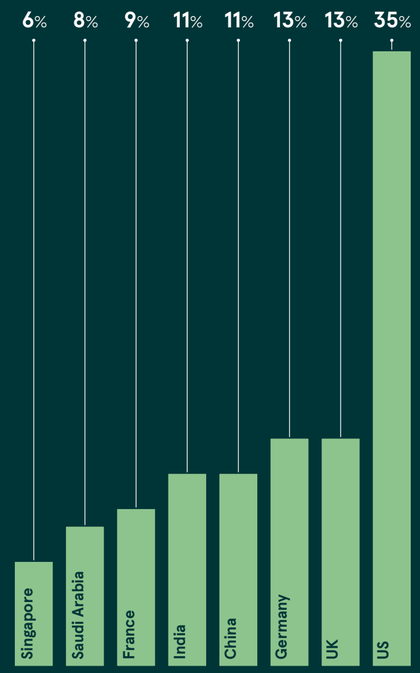

PwC, 2026

54%

of enterprises expect to have at least 40% of their AI experiments in full production within the next six months, a massive jump from just 25% today

PwC, 2026

91%

of customer service leaders report intense executive pressure to implement AI in 2026, leading to 80% of organizations planning to redefine human agent roles

Gartner, 2026

65%

of CMOs believe advances in AI will dramatically transform their role within the next two years, moving from influencers to 'designers of business impact'

PwC, 2026

60%

of the global workforce now has sanctioned access to AI tools, a significant increase from under 40% just one year ago

OpenAI

ENTERPRISE PRIORITIES AND MARKET RISKS IN 2026

Ranked by percentage of leaders citing as a high/top priority or significant concern

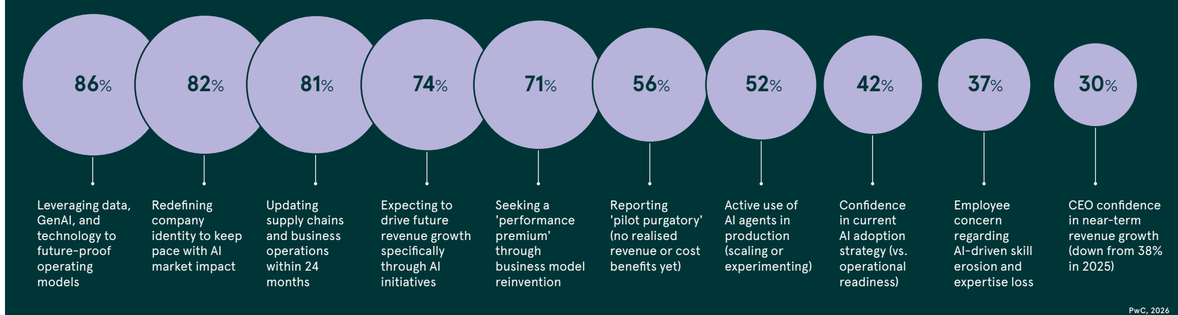

PwC, 2026

AUTOMATION

The autonomous enterprise: how the C-suite is becoming agentic

As AI moves from conversation to execution, a new wave of autonomous tools is rebuilding business operations from the plumbing up

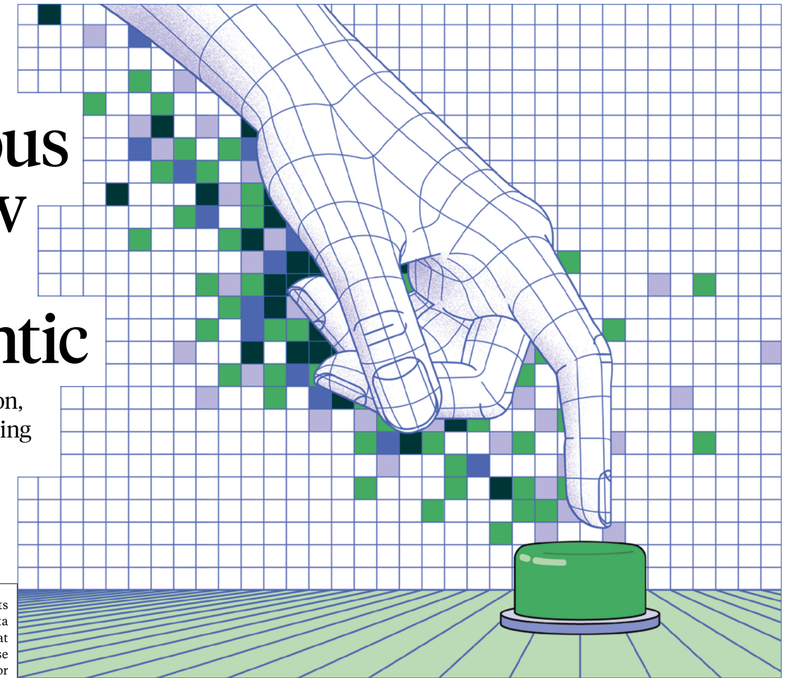

Tom Dennis

The initial excitement over generative AI is cooling into a more practical, and arguably more powerful, reality for the C-suite. For years, the singularity has been discussed as a distant, theoretical horizon where machines surpass human intelligence. But for the modern business leader, this is a distraction. The real story of 2026 is not a future prophecy; it is the immediate rise of 'agentic' AI – autonomous systems that do not just talk, but act.

While early experiments focused on surface-level search and simple chatbots, the next stage of transformation is defined by enterprise agents. These are autonomous systems that carry out complex, multi-step workflows without constant human intervention. For leaders navigating this transition, the focus is shifting from what AI can say to what it can do. This evolution is being driven by new tools that connect intelligence directly to firms' operational plumbing.

Anthropic's 2026 Agentic Coding Trends Report shows a meaningful inflection point in adoption. Over half (57%) of organisations now deploy agents for multi-stage workflows, while 16% have progressed to cross-functional processes that span multiple teams. This transition suggests that AI is no longer a peripheral experiment, but a core piece of infrastructure.

Anthropic is leading much of this shift by positioning its Claude model as a 'next-generation colleague' for the office. By prioritising safety and factual accuracy over creative flair, the firm is moving away from general consumer assistants. Instead, it offers tools that can ingest an entire codebase or hundreds of legal documents to provide strategic value.

For leaders, this represents a move from human-in-the-loop to human-on-the-loop. Instead of micro-managing every prompt, executives are now overseeing autonomous agents that handle the heavy lifting of data synthesis. Nicola Johnson, CFO at Pulse Clean Energy, notes that these tools are becoming essential for focus. "It just removes the clutter and cuts down the admin more than anything," she says. "We want to be as efficient as possible with the bits that aren't value-adding, so we can focus on the strategic bit."

The greatest barrier to implementing these systems is not intelligence, but integration. Many AI tools remain trapped as expensive dashboards, largely because they lack real-time access to operational data. To be useful, an agent must be able to see the current state of the business and do something about it.

Traditional systems rely on slow, polling-based architectures where data is checked at intervals. Services like HiveMQ and EMQX replaces this with event-driven streaming. Using a lightweight IoT protocol these services stream real-time data from devices to AI models. It acts as a fast, reliable backbone for feeding data to AI models for predictive maintenance, anomaly detection, and autonomous systems This allows AI agents to observe and act on telemetry – such as power usage or production quality – the moment it becomes relevant. When an agent can react to a sensor in milliseconds, it stops being a report-writer and starts being an operator.

Similarly, developers like Axoniq and JellyFish as well as AWS and Azure are enabling agents to interact with complex internal systems through the Model Context Protocol (MCP). This standard allows an AI agent to send a natural language prompt that maps directly to a system command. It ensures business logic runs just as it would for a human employee, but with the speed and scale of software.

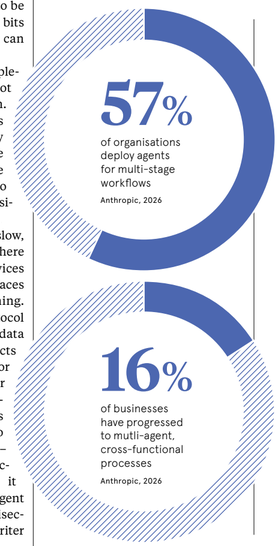

> **To be useful, an agent must be able to 'see' the current state of the business and 'do' something about it**

Rachita Sundar, CFO at Qualtrics, believes the transition requires a change in mindset from the top down. "The skills I built as a software engineer, such as critical thinking, curiosity and problem-solving, serve me very well in business," she says. "I was used to looking at everything as white space, as an opportunity." For the modern C-suite, the 'white space' is now the gap between data and action that agents are beginning to fill.

In the business and marketing space, AgentPress is showing how these tools transform business case planning and sales enablement. Rather than using AI to merely write drafts, firms use agents to orchestrate custom business cases in seconds, researching accounts, use cases, products, and ICPs automatically in a single, automated loop.

This shift is forcing a total rethink of performance and talent. When an agent can handle the SEO strategy or the initial financial audit, the value of human staff shifts toward oversight and ethics. Pearson, CHRO Ali Bebo is rebuilding performance management to keep pace with these shifts. She argues that traditional long-term planning is now outdated in a time of swift AI advancements.

"Traditional long-term planning, anchored in annual or even 18-month cycles, is no longer fit for purpose," Bebo says. Instead, she advocates for a more agile approach where leaders "act their way into the right thinking" rather than trying to plan every outcome in advance. Her 'GPS' model for skills allows the organisation to recalibrate in six-month increments, ensuring that human talent remains complementary to machine capability.

The move toward an autonomous enterprise requires more than just a tech budget; it requires a new leadership philosophy. As AI agents begin to handle the 'clutter' of the modern office, the role of the human leader becomes more about setting the guardrails and the vision.

However, the cost of standing still is higher. The impact of AI is already visible in workforce changes throughout 2025. Many companies now use AI tools to boost employee productivity while simultaneously reducing recruitment for junior roles. Leaders must maintain a sufficient pipeline of talent to avoid long-term capability gaps, even as they automate entry-level tasks.

For the board, the lesson is clear: the singularity is a distraction, but the agent is a reality. Success in 2026 depends on how well a company can integrate these autonomous tools into its core processes. This requires a focus on 'data readiness' and a willingness to dismantle the silos that prevent agents from accessing the information they need.

As Deputy's CFO, Emma Seymour, notes: "AI agents won't just be tools, they'll be teammates." This means treating AI implementation not as a software installation, but as a change management project.

The autonomous office is a practical reality that demands a response today. Leaders who focus on the plumbing will be the ones who scale successfully. Those who wait for a mythical singularity will find themselves outpaced by the agents already working next door.

Commercial feature

Securing AI: the identity challenge behind the AI boom

As AI agents scale rapidly, organisations must address identity, governance and control to ensure secure, accountable and trustworthy deployment

Businesses are racing to adopt AI, drawn by the promise of total transformation. Over the past year, AI agents have emerged as a new digital workforce – analysing transactions, triaging customer queries, assisting clinicians and even making operational decisions.

But this rapid deployment creates a new challenge for leaders: identity. As employees experiment with tools outside formal governance, firms face a wave of 'shadow AI'.

In this environment, identity is hard to manage. It is difficult to track how agents interact with systems, what data they access and whether they comply with internal policies.

Indeed, research from Okta shows that 91% of organisations are already using AI in some form. Yet only 10% have a well-developed strategy for managing, and importantly securing, non-human identities.

In other words, while businesses are racing to deploy AI, very few are equipped to control it.

While AI adoption often starts small – a pilot project or proof of concept – things can quickly escalate.

"Typically, organisations begin with simple, measurable tasks. Then very quickly, they want to scale. They want to go bigger and faster," says Stephen McDermid, CSO EMEA at Okta.

That's when AI agents move from isolated tools to embedded operators, interacting with systems, accessing sensitive data and, increasingly, acting autonomously. It's also when governance often falls behind.

Without clear controls and governance, organisations can quickly lose control of how their AI agents behave – what they have access to, how they interact with systems and whether their actions align with policy. According to Okta, only just over half of organisations have full visibility into their AI activity, leaving the rest exposed to significant blind spots.

AI agents need to be treated as first-class identities, meaning they should be identified, authenticated and governed in the same way as human users. They log into systems, they access applications, they retrieve and share data. In many ways, they behave like employees – but currently without the same guardrails.

> **While businesses are racing to deploy AI, very few are equipped to control it**

"When you look at how these agents operate, they're acting on behalf of users, customers or services. So they need to be treated as first-class identities in their own right," explains McDermid.

Yet in many organisations, they're not. AI agents are deployed without robust authentication, clear permissions or lifecycle governance, which raises critical questions as AI agents take on more responsibility: who is in charge? They might approve a transaction, respond to a customer or access a database. But who authorised that action? What level of access was granted? And who is accountable if something goes wrong?

There have already been cases where poorly secured AI systems have exposed sensitive data or been manipulated into unintended behaviour. In one instance, a recruitment chatbot leaked millions of records due to weak credentials. In others, attackers have exploited AI systems through prompt injection to bypass controls.

"The value of AI comes from the data it can access. But as you connect more systems and add more data, you can very quickly lose control," says McDermid.

This increasing gap between what AI can do and what organisations can control is what security leaders are calling the "authority gap". To close the gap, organisations need an approach that treats AI agents as first-class identities. This is where identity platforms come into play.

Okta's unified identity platform is designed to bring order to this complexity, providing a central layer through which all identities – human and non-human – can be managed, authenticated and governed, underpinning its blueprint for the secure agentic enterprise.

"Identity gives you the ability to control access, monitor behaviour and enforce policy. It becomes the foundation for trusted AI," says McDermid.

In practice, this means authenticating every agent interaction to ensure only approved systems can operate, and controlling access at a granular level so agents only see what they need to see. It also involves monitoring behaviour continuously to identify anomalies or misuse in real time, and governing the full lifecycle from creation through to decommissioning.

Identity platforms such as Okta's enable a shift from reactive security to proactive control, allowing organisations to scale AI safely.

The good news for business leaders is that governance and innovation are not mutually exclusive.

Companies such as Siemens are embedding identity-led security into their digital transformation efforts, ensuring that every user, system and agent is governed consistently. McLaren Racing, operating in a high-performance, data-intensive environment, has made identity central to both security and operational agility.

In financial services, firms like Paysafe and Equals Money are integrating identity controls directly into AI-enabled products, aligning innovation with regulatory expectations from day one.

The common thread among those businesses is that identity is not an afterthought. It is built in from the start.

Okta, 2025

There are some practical first steps that organisations that are still early in their AI journey should consider.

- First, find your shadow AI. Assume AI is already in use, whether sanctioned or not. Identify where it exists, how it is being used and what data it touches.
- Second, set the rules of engagement. Define clear policies for AI use and educate employees. Most security risks still originate with people – and AI is no exception.
- Third, put visibility and control in place. Deploy tools that allow you to monitor, govern and manage AI agents as they evolve. Without visibility, there is no control.

These steps are not about slowing progress, but about enabling it, safely and sustainably.

AI is reshaping how organisations operate. But as it becomes more autonomous, the stakes are rising. The question is no longer whether to adopt AI, but how to do so responsibly.

McDermid notes: "To get AI security right, you have to get identity right."

Those that invest in identity – establishing clear governance, visibility and control – will be best positioned to unlock AI's full potential. Those that don't risk falling into the authority gap.

For more information please visit okta.com/solutions/secure-ai

okta

DIGITAL SOVEREIGNTY

Breaking the cloud vendor trap

Rising costs and geopolitical shifts are ending the era of total cloud reliance. Here is how business leaders are clawing back control of their data and infrastructure

Joy Persaud

Cloud computing was once the natural option for enterprises, shaping the way systems are structured and run, shifting workloads to hyperscale platforms. However, issues around rising costs, geopolitical uncertainty, and data residency regulations are forcing tech leaders to claw back control.

Cloud repatriation from dominant vendors like AWS, Google Cloud and Microsoft can return operational and jurisdictional control to organisations. It also helps businesses avoid hidden costs within the 'Big Three' vendor trap. Systems are easy to build on these platforms, but they often rely on proprietary databases, AI services and analytics tools that suffer from unpredictable costs and outages.

James Lovegrove, public policy director (EMEA & APAC), Red Hat, says EMEA leaders are increasingly focusing on digital sovereignty, where the legal and technical capacity to audit, modify and secure one's own environment, according to regulatory requirements, is retained.

"EMEA leaders are moving beyond the cloud efficiency drive to a compliance and resilience mandate, both in response to regulatory and wider policy trends and as part of their business logic," he says.

"Over-reliance on a single proprietary provider can undermine an enterprise's ability to compete as well as comply with EU regulations or qualify for future procurement opportunities in the region. Location of data is a risk factor, but more importantly organisations are seeking to preserve operational autonomy."

The move towards digital sovereignty is already ramping up. Gartner forecasts that sovereign cloud infrastructure as a service (IaaS) spending will shift 20% of current workloads from global to local cloud providers. Also, global sovereign cloud IaaS spending is expected to hit $80bn this year, representing growth of 35.6% year-on-year. In Europe alone, growth could increase by 83% in the same period, Gartner finds.

Clearly, digital sovereignty is driving major shifts in enterprise structure investment. But is sovereignty cost-effective? Analysis carried out in 2025 by BCG Global indicates that sovereign cloud options come with premiums of 10-30% compared with public cloud offerings. This is because of compliance controls, isolated infrastructure and region-specific staffing. Therefore, sovereignty is likely to increase short-term infrastructure costs while reducing long-term strategic exposure.

When the cost of leaving the cloud becomes too high, the benefits of hyperscale diminish. This 'lock-in threshold' is the point at which the cost of moving data, modifying applications and reskilling employees exceeds the savings that cloud services promised.

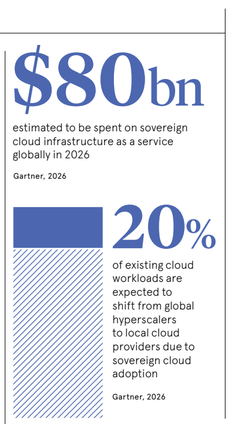

> **The real power of AI comes from aggregating data at scale. The future isn't edge or cloud, it is both**

Provider-specific tools such as analytics, AI services, managed databases and frequently accessed datasets make migration pricier yet. And, the more embedded in the cloud an enterprise is, the more complex the problem. Transferring data can cause operational disruption, delays, and involve hefty fees.

To retain flexibility and keep reaping the benefits of hyperscale and cloud innovation, CIOs and CFOs must be clear about this tipping point and recognise just how locked in they are. They can then choose which data-heavy, more predictable, latency-sensitive systems can be shunted to edge or private systems where costs are easier to control.

Emma Lauchlan, director of growth, Asanti Data Centres, says: "Leaders are drawing a clearer line between workloads that can move for efficiency and those that need to remain in tightly governed environments to meet GDPR, sector regulation and customer expectations around privacy and trust. It is one reason why a third of organisations plan to move more workloads into on-premises or UK colocation environments, so they can retain a higher degree of control over where data resides and how it is protected."

When it comes to cost savings, cloud repatriation is often positioned as a move from OpEx to CapEx but the reality is not cut and dried.

Dean Garvey-North, CTO of Microlise, explains: "Cloud shifts the investment focus toward innovation rather than infrastructure management. However, edge should be viewed as complementary to hyperscale cloud rather than a replacement. CapEx typically increases due to hardware procurement, facilities, and infrastructure lifecycle management. Operational complexity and staffing requirements rise, often offsetting perceived OpEx savings. The ability to scale elastically and innovate quickly can be reduced."

The true value for most companies lies not in owning infrastructure but in the software, platforms, and data they build on top of it, he adds, advocating a hybrid model whereby edge is in place for real-time operational intelligence and cloud is deployed for data platforms, AI training and large-scale analytics.

But, what exactly comprises a sovereign tech stack in the current regulatory environment? Control of the full dependency chain is key. Wayne Scott, GRC solutions lead at Escode, says sovereignty depends on independent access to critical applications regardless of where the supplier is headquartered, clear legal rights across the software supply chain, the ability to validate and rebuild software if a vendor fails, and transparency over where applications sit and who ultimately governs them.

"Also, it's not just about data residency anymore," he says. "You can have data sitting in the right jurisdiction, but still be completely dependent on a supplier you can't exit or replace. Regulators are starting to look beyond location and ask whether firms can actually continue operating if that supplier fails, withdraws or is no longer viable."

It is no longer an 'if' but rather a 'how' enterprises achieve that hybrid sweet spot between hyperscale and edge that limits risk and dependency, and maximises innovation. A new, resilient infrastructure will live or die on the ability of CIOs and CFOs to assess long-term value coupled with risk, where factors such as cloud cost, the impact of outages, price rises, and the difficulty of moving workloads fall under their control.

As Garvey-North puts it: "Edge computing brings intelligence closer to operations, but the real power of AI comes from aggregating data at scale. The future isn't edge or cloud, it's an architecture that intelligently combines both."

THE RACONTEUR 50

Recognising those who lead.

The role of the modern-day CEO is evolving. It is no longer enough to focus solely on profit, revenue or share price. Leaders must balance financial performance with employee wellbeing and ESG concerns, finding ways to innovate and grow at a time of deep uncertainty and turmoil.

Across five categories, we hope that by shining a spotlight on the best business leaders, we can offer insights into what it takes to lead from the top and inspire the CEOs of the future.

Meet the 50 CEOs changing British business.

raconteur.net/
the-raconteur-50-2025

Raconteur