INDEPENDENT PUBLICATION BY

Raconteur

RACONTEUR.NET | #0937 | 19/06/2025

ENTERPRISE AI

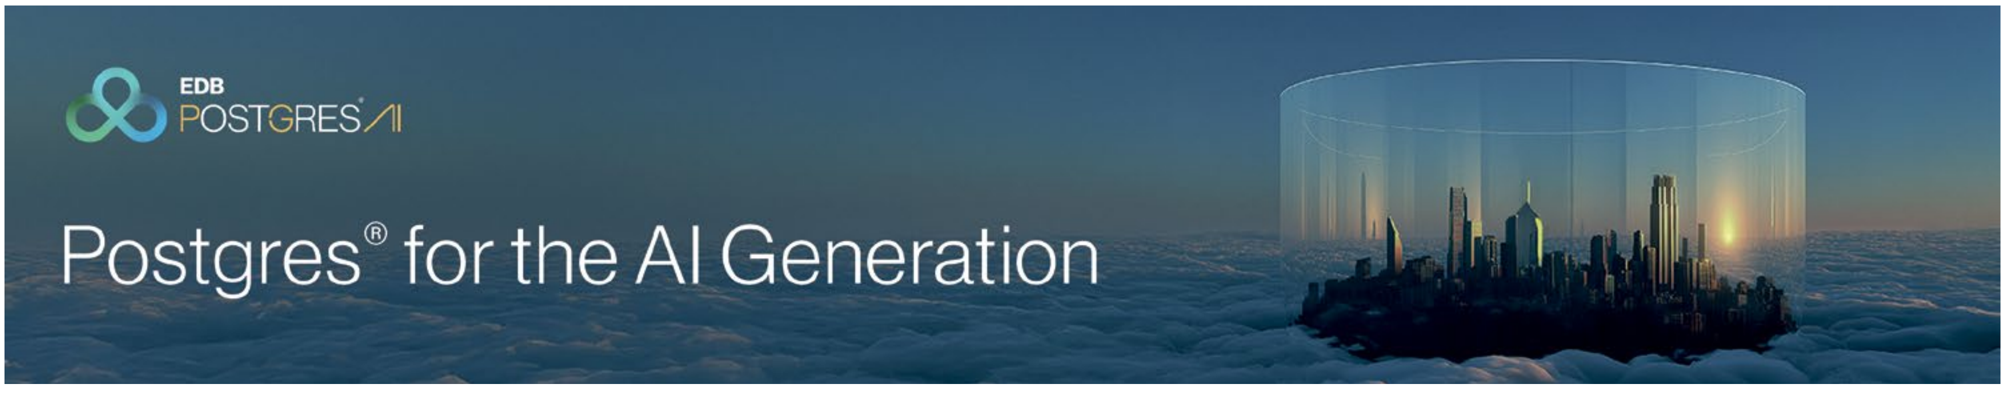

ENTERPRISE AI

Distributed in
THE TIMES

Contributors

Sam Forsdick
Raconteur's HR editor. He analyses the trends shaping the function, offering insights and advice for people professionals.

Tamlin Magee
Senior technology writer at Raconteur. He's interested in big ideas shaping business tech and the impact of new technologies on people and society.

Raconteur

Special projects editor
Ian Deering

Commercial content editors
Laura Bithell
Larnie Hur
Jessica Lynn

Commercial production managers
Alex Datcu
Ellen Newsome

Production executive
Sabrina Severino

Design and illustration
Kellie Jerrard
James Lampard
Celina Lucey
Samuele Motta

Design director
Tim Whitlock

Certified B Corporation

Although this publication is funded through advertising and sponsorship, all editorial is without bias and sponsored features are clearly labelled. For an upcoming schedule, partnership inquiries or feedback, please call +44 (0)20 3877 3800 or email info@raconteur.net.

Raconteur is a leading business media organisation and the 2022 PPA Business Media Brand of the Year. Our articles cover a wide range of topics, including technology, leadership, sustainability, workplace, marketing, supply chain and finance. Raconteur special reports are published exclusively in *The Times* and *The Sunday Times* as well as online at raconteur.net, where you can also find our wider journalism and sign up for our newsletters.

in raconteur-media   @raconteur.stories

raconteur.net /enterprise-ai-june-2025

GENERATIVE AI

Can tech users opt out of GenAI?

Generative AI has become standard in our digital lives, with many apps now sporting AI features. So can ordinary tech users avoid interacting with AI if they want to?

Tamlin Magee

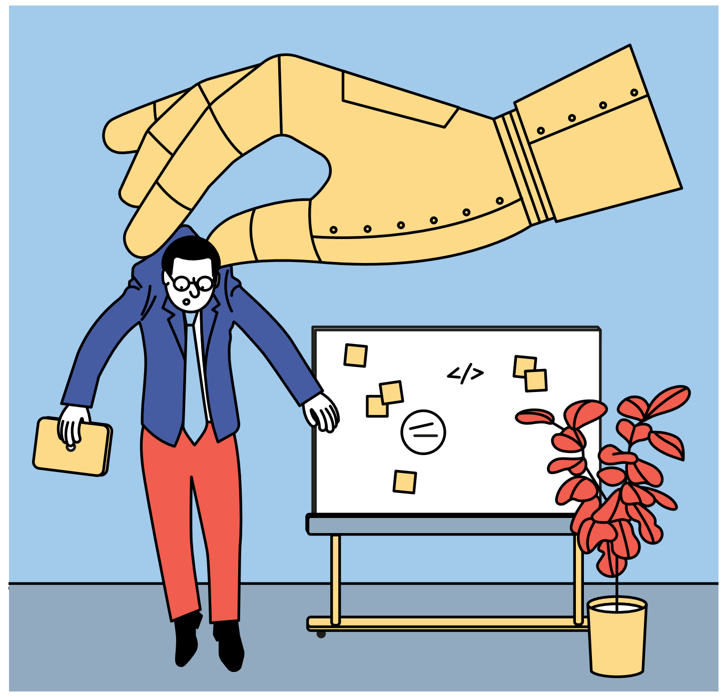

Nearly every major digital platform markets AI features. In our professional lives, AI pops up to summarise our Zoom meetings or help us to navigate the Salesforce CRM. AI is also there to offer assistance on personal apps such as Instagram and WhatsApp – often to the annoyance and frustration of users, however.

Companies will no doubt expand their use of AI, too. Providers have convinced enterprise users that the technology will bring significant increases in productivity and revenue. Most companies, however, have failed to achieve any bottom-line results from AI adoption, according to *The State of AI* report by McKinsey.

The widespread use of these tools, and the prevalence of AI-generated content, invites questions about the ownership and control of our personal data. Individuals can choose not to interact with platforms such as ChatGPT and companies can ban staff from using the technology – a quarter of businesses have taken this step – but it is increasingly difficult to avoid AI altogether.

"How do you opt out of a technology that is used by others to interact with you?" asks Bruce Schneier, a public-interest technologist and chief of security architecture at Inrupt, a decentralised data platform. "Let's say it's used to write fundraising emails from political credits. You can't tell your email programme to delete all messages written by AI."

He likens it to opting out of handguns: "You can choose not to buy one, but you can't choose not to get shot by one."

GenAI is intrusive by design. These systems rely on data to train the large language models that power them. And where does the data come from? Long-forgotten forum posts, half-baked opinions posted on social media, real conversations between friends or strangers, artwork, poems, blog posts and so on. Any content on the public internet is up for grabs.

Sometimes AI systems are fed data outside of the public domain, too. For instance, Meta allegedly used material pirated from a digital book repository to train its models. And, in pressing ahead with its opt-out approach to copyright protections for art and AI-generated content, the UK government is in practice placing once-protected content in the public domain by default.

According to Carissa Véliz, an associate professor of philosophy and ethics at the University of Oxford, it may be impossible to opt out of GenAI; that is, to stop it from collecting your personal data. That's a problem, she says, because it means we have essentially lost our basic privacy rights and those protected by the GDPR. "These companies don't even know what data they're using," Véliz says.

Even if companies did track the data they use more effectively, they would be unlikely to delete it in the face of public pressure, she claims. Doing so would mean scrapping their models and re-training them without any contested data. "That's not going to happen," Véliz says.

Jaya Klara Brekke, chief strategy officer at Nym, a privacy-technology company, says it is almost impossible for individuals to opt out of GenAI in any meaningful way. "The nature of AI and LLMs is that they work in the aggregate, meaning even if you manage to opt out of training data, someone else won't," she explains. "The collection of everyone else's data then sets a norm."

There may however be a technical solution to this issue, according to Matthew Hodgson, co-founder and CEO of Element, a private messaging application. By adopting a mechanism to track and separate AI-produced content from material made by humans, society could set a precedent for defining the integrity of information, he explains.

Hodgson imagines a future where individuals own any data they produce, can track where it goes and can prove that it was made by them. They could then specify whether or not they want that data to be used to train LLMs. If your data comes up in a model and you haven't given consent for its use, you can prove that it was stolen and used against your wishes, he explains.

"Regulation is just not enough to protect the digital rights of people, including privacy rights, creative rights and many others that the AI industry regularly violates," adds Klara Brekke.

Digital rights, she says, must be part of the infrastructure by default to be effective. But, more importantly, alternatives must be "supported, funded and developed. It has to be made clear that there actually are alternatives."

Klara Brekke says the regulation of AI and the technologies supporting it must develop in tandem with investments in the industry. "The EU is starting to wake up to the fact that big tech can't just be regulated away," she says, "but we need that investment in alternatives."

So GenAI is here to stay. And if we can't opt out of it, perhaps the tech companies that create it should be held more accountable for its misuse through regulation and stronger data ownership rights. ●

AI MODELS LEARN FROM EVERYDAY INTERACTIONS

Approximate monthly web visits to GenAI platforms

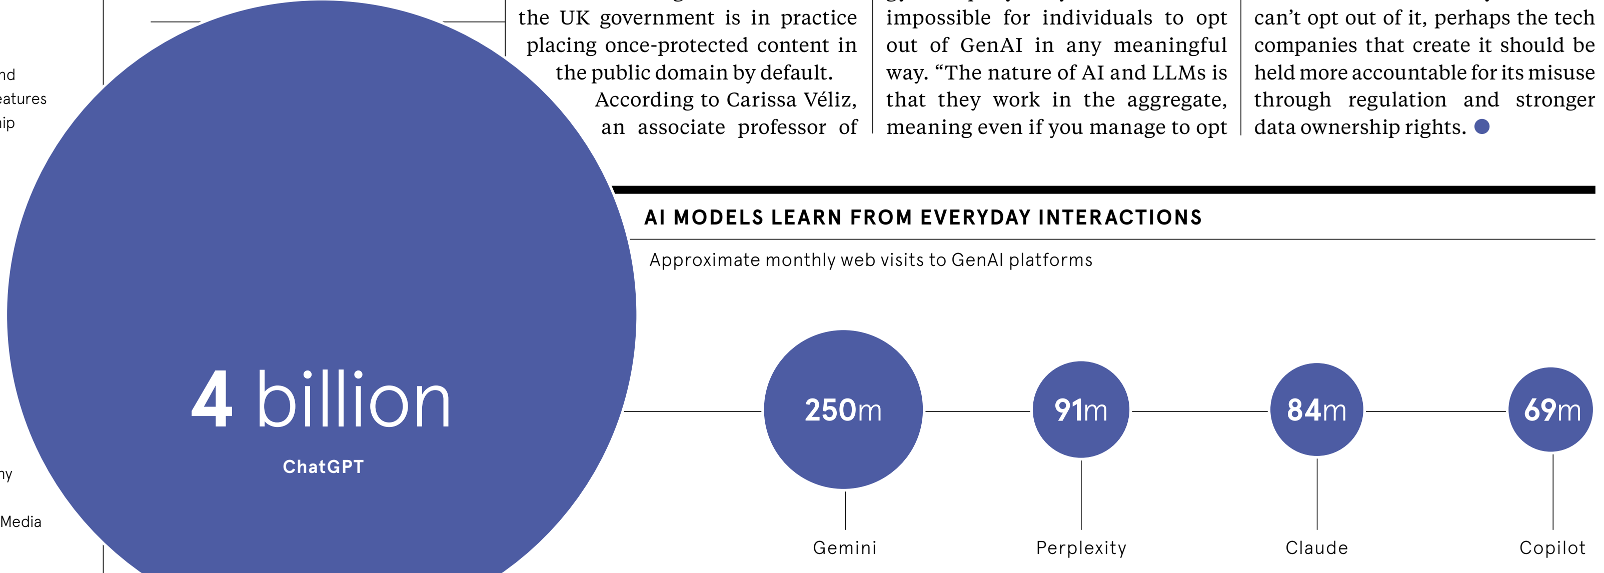

Providers' websites, 2024 to 2025

Commercial feature

From talk to transformation: rethinking AI to unlock human potential

Making AI implementations more human-centric could radically improve how people work, collaborate and create

The way we work isn't working. UK employees are losing nearly two full working days each week to low- or no-value tasks – from form-filling to chasing approvals – with nearly 45% of administrative work considered inefficient, according to PwC's UK Workforce Hopes and Fears Survey 2024.

AI promised to improve this situation by taking on the tasks that stop people from doing engaging and impactful work. But as many organisations are discovering, it's not quite that simple.

Layering AI onto flawed processes and workflows rarely results in long-term productivity gains, let alone more time for high-level work. In fact, it can actually make life more complicated for teams.

That's not to say AI isn't a potent tool for knowledge workers: research by Miro found that 76% believe AI could benefit their role.

However, leaders need to rethink their approach to implementing the technology to ensure their teams – and the organisation – truly reap the benefits.

Today, for instance, 54% of workers struggle to know when to use AI, while 35% describe their AI skills as "nonexistent".

They are also getting mixed signals from their organisations, which adds to the confusion around the technology.

For example, 39% report that their company often abandons AI efforts, while 46% agree that there is more talk than action.

"The complexity of bringing AI to organisations is often due to thinking it is just a [technology] implementation," says Tomás Dostal Freire, CIO and head of business transformation at Miro, a collaboration platform with AI features designed to streamline workflows and accelerate innovation.

"You need to take a step back and really think through how AI can transform how you operate."

Amplifying human skills

The hype around AI is part of the problem. All those articles and LinkedIn posts promising mind-blowing productivity gains with one simple tool have arguably created unrealistic expectations about what can be achieved.

But the thing leaders really need to think about is: "What does it mean to the organisation – not in terms of full automation and replacing people, but rather as an augmenting force for employees?" asks Dostal Freire.

This demands a different kind of leadership approach – one that asks not how much AI can automate, but where it can amplify people's ability to collaborate, innovate and solve complex problems.

Smarter implementation of AI focuses on where teams tend to lose momentum, for example – perhaps due to information silos, excessive time spent searching for resources, or an inability to ideate at speed – and how the technology can address these friction points.

A human-centric, rather than tool-led, strategy focuses on how AI can help people connect and contribute more effectively, strengthen feedback loops, and innovate at speed.

Organisations that successfully leverage AI to amplify human potential tend to focus first on the desired outcomes, then, in the following order, the people, processes and technology needed to achieve them. Crucially, they place a strong emphasis on employee education.

Indeed, employees should be encouraged to upskill not only for the employer's benefit, but for their own long-term career value. For as Dostal Freire says, "AI literacy is the new digital literacy."

LASTING VALUE COMES WHEN DECISION-MAKERS TREAT AI AS A PARTNER THAT GOES BEYOND EFFICIENCY TO UNLOCK TRUE HUMAN POTENTIAL

Employees citing the most significant benefits of AI adoption

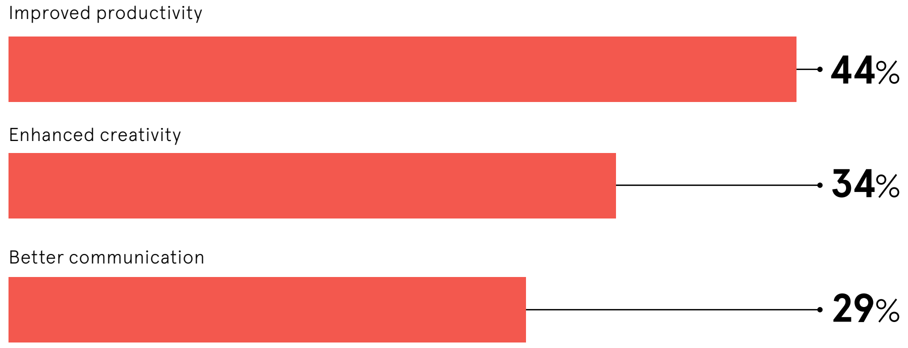

Miro, 2025

> Organisations should focus on three interconnected principles: educate, inspire and empower

Inspiring employees

Despite this fact, initiatives that support and develop AI literacy are clearly lacking today. Notably, formal training is the number one thing the 8,000 global employees recently surveyed by Miro said would help them feel more confident about adopting AI.

This training needs to reflect the nuances of AI usage. "When you think of a traditional technology implementation, you usually end up training people in how to use the tool, so it's a tool-centric approach," says Dostal Freire.

"We need to shift towards a people-centric and workflow-centric approach. So it's no longer only about 'how to' use AI, but also a lot more about 'when to' and 'why to' use AI."

To successfully implement human-centric AI, organisations should also focus on three interconnected principles: educate, inspire and empower.

The education element, as mentioned, moves beyond tool training to help people understand context and application.

Inspiration, meanwhile, involves showing relevant examples of how it's successfully amplifying human ability.

"Show the art of the possible. Even if you think you know what AI can do, you still need to share what is best-in-class from peers…and it should not be AI at large, but for a team or department – so finance, marketing, etc." Dostal Freire explains.

Finally, empowerment means creating safe environments for employees to use the technology to its full potential. "Once they know what AI is, and they're pumped about what it can do…you need to give them platforms to play around with it and discover it for themselves."

This requires security teams to become enablers rather than blockers, adapting governance at the speed of AI innovation—a challenge that can be met through close collaboration between the CIO, CISO, and other members of the security leadership team.

Measuring success

The way leaders gauge the success of AI implementations also requires a rethink. Measuring how many people are using a platform, for example, is much less relevant a metric than whether the business outcomes that were the driver for adoption have been achieved.

The concept of Return on Employee (ROE) can also complement ROI when evaluating AI's impact.

Rather than simply focusing on the financial return, it reveals the broader value of an implementation by examining how it impacts employees, factoring in things like job satisfaction, improvement in collaboration and the quality of work produced.

"It's harder to quantify…but if you have more engaged employees and faster decision-making, ultimately you do see the results in better outcomes," says Dostal Freire.

Almost two-thirds of workers agree that AI can improve wellbeing and job satisfaction, for example, which, in turn, can bolster productivity and innovation. More than a third also believe it can enhance creativity, while 29% believe it can lead to better communication.

As AI advances toward more sophisticated applications, including autonomous agents, the human amplification model will become even more critical for success.

Dostal Freire says that empathy should inform AI deployment decisions and is a powerful way to assess which processes should be automated and which should remain primarily human-led.

High complexity, high empathy activities would involve a human taking the lead, for example, perhaps with AI as a copilot.

"However, where there's low human touch advantage and high complexity, or low human touch advantage and high repetition, that's where you could rethink and automate," says Dostal Freire.

In the end, the best path forward isn't about choosing between human capability and artificial intelligence. Instead, it's about creating conditions where AI amplifies what makes humans most valuable: the ability to collaborate, innovate and solve complex problems together.

Commercial feature

Is secure AI the smartest tool in a leader's innovation kit?

Businesses must embrace the opportunities that AI provides to push the boundaries of innovation, while ensuring sensitive data is kept safe

As AI adoption accelerates across every sector, the pressure is on for organisations to harness its potential – from boosting productivity to cutting costs and driving innovation.

But these benefits risk being lost without a clear, secure AI strategy in place. Without proper safeguards, businesses could expose sensitive data and erode customer trust just as quickly as they scale.

The real challenge is finding the balance: how can organisations unlock the power of AI while staying in control of their data?

The answer lies in robust strategy and governance. With the right foundations, AI tools can become a genuine competitive advantage – helping teams work smarter, save time, and innovate with confidence.

"Business leaders across the world are grappling with how to use AI tools effectively and safely to speed up processes and improve their organisations," says OneAdvanced's chief product officer Amanda Grant.

"Relying on tools that learn from the information uploaded is too risky. Organisations need confidence that their data remains safe."

AI is being used increasingly across all sectors of the economy, including in healthcare, housing and finance.

In healthcare, it is reshaping how care is delivered and decisions are made by enhancing diagnostic accuracy, enabling more efficient administrative workflows and improving patient engagement.

In hospitals, for example, AI systems are being employed to forecast hospital resource demand, allocate beds and equipment more efficiently and to optimise supply chains.

Away from hospitals, Mastercall Healthcare, a social enterprise which provides outpatient treatments, has looked to OneAdvanced AI to reach its business goals and become more productive.

of organisations identify their own employees as their main point of cyber risk

OneAdvanced, 2025

Their Large Language Model (LLM) is designed for business use and is tailored around a specific organisation and sector. Automation is embedded into the workflows within the software platform to provide a fully encrypted, private and closed service.

This is important because it addresses the growing concern around shadow IT.

A survey of 6,000 knowledge workers by Software AG reveals that half of those asked were using non-company-issued AI tools, and 46% admit they would refuse to give them up even if it were banned by their bosses.

Mastercall Healthcare's director of digital innovation, Jonathan Ritchie, says implementing private, secure AI is now a critical necessity.

"This approach enables us to leverage AI's capabilities while maintaining strict control over our sensitive data," he says. "We enhance our security posture but gain the ability to customise our AI environment. This drives innovation without compromising on compliance."

Indeed, the most effective and resilient innovative AI strategies are built on a foundation of secure data, and organisations need to know how and where their data is being processed, especially in mission-critical sectors where safeguards are vital to ensure personal data is not shared outside of their own organisation.

Commercial feature

LEADERS WHO WANT TO DRIVE INNOVATION AND UNLOCK AI'S FULL POTENTIAL MUST CONSIDER HOW TO CREATE A RESILIENT AND SECURE DATA FOUNDATION

Percentage of leaders citing their key areas of focus for genAI risk mitigation efforts

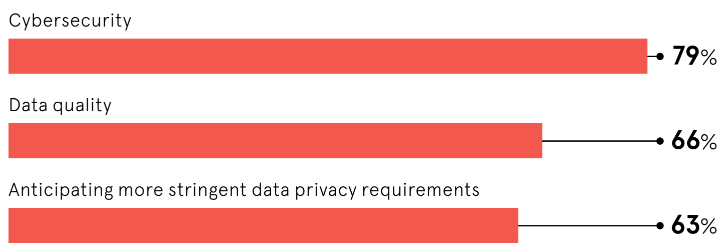

KPMG, 2024

In fact, companies can face large fines if they claim customer data is being processed in one territory, but in reality it is being handled somewhere else.

The new OneAdvanced AI software platform offers complete UK data sovereignty.

"If you do not have controls on how your data is used, it can be used in ways that weren't intended. You could be breaking the law as well as putting your business at risk," says Grant. "You need to know where your data is going when you switch on an AI tool. With our platform your data is fully processed, backed up and recovered in the UK."

Organisations often worry about balancing innovation at pace with security. There is clear evidence that when there is a robust AI strategy in place and tools are integrated effectively, there can be massive productivity gains.

Many organisations are already benefiting from AI's ability to solve complex problems using their business's own data. This can be central to innovation around product development or finding new ways to raise customer revenues.

Grant accepts that for most organisations, implementing AI will be a stepped journey as they balance real and perceived security risks.

She urges companies to experiment with AI to devise or trigger innovative ideas or concepts that will benefit them and the business by solving current challenges.

Her advice is not to focus just on the short term and the next step, but to think about what full autonomy could mean for your business and the workstreams that it can transform.

"It's a necessity to build the foundations today to ensure that you keep pace with the shifting landscape," she says.

"The first step of the journey is to learn what AI can do for your organisation. Often, there needs to be a shift in mindset and an education of employees, especially as the pace of change in AI speeds up."

Grant suggests creating a no-risk playground where people have the freedom to try innovative ideas with AI to familiarise themselves with the technology and learn how it can make the business more successful and their jobs easier. This might be a simple task, such as using AI to summarise a meeting.

> "By coupling innovation with robust security standards, it is possible to drive economic growth while protecting sensitive data

AI assistants can free people from admin-heavy tasks, such as comparing a procurement policy with supplier contracts to check compliance.

It can also improve efficiency using Retrieval-Augmented Generation. This is where the technology accesses organisation-specific documents and data to tailor responses from the LLM.

It can connect to external sector-specific data sources, such as legal databases, to understand and answer questions. An AI tool could then be used to help draft a tailored response to a complaint by crafting a reply around the company's policies or procedures.

OneAdvanced's CEO Simon Walsh sums up the dual challenge facing organisations today: managing a rapidly evolving technology landscape while driving strategic growth and maintaining operational efficiency.

"By coupling innovation with robust security standards, it is possible to drive economic growth while protecting sensitive data," he says.

"Our OneAdvanced AI platform represents a crucial step forward in ensuring that organisations can seize these opportunities without risk."

His message is clear – in an AI-driven future, progress and protection must go hand in hand.

For more information please visit oneadvanced.com

oneAdvanced

Q&A

How leaders can build secure strategies

How can businesses embrace AI confidently and securely? OneAdvanced's chief product officer **Amanda Grant** shares practical steps to drive innovation while safeguarding data and engaging employees

Q How easy is it to understand AI's strengths and opportunities and to be secure?

A Secure-by-design AI must be a key player in cultivating innovation. You need to identify areas where AI can automate tasks, provide insights and enhance user experiences to accelerate your business's 'jobs to be done'.

For successful AI adoption, employees need to clearly see how it adds value to their work. While younger team members may adapt to AI seamlessly, older staff might need extra support and guidance to understand how it can simplify their responsibilities and boost productivity. With the right approach, AI can become a valuable ally for every member of your team.

Q When it comes to using AI and defining business objectives, how should the two work together?

A It is important to communicate effectively how AI will help a business achieve its goals. The first task is to define what those goals are by analysing what problems the organisation needs to solve, and then to work out which of these opportunities could be capitalised on if AI were used. Of course, you must ensure you have robust metrics in place to measure success.

Set up a governance framework, including a steering committee, guidelines and user training to ensure effective AI adoption and use. This will help employees to understand, for example, the risk of using shadow IT as well as appreciate how AI can help in their role and the wider business.

Q Securing data and using data effectively to fuel innovation is crucial here, so how do we get this right?

A Firstly, place safeguards to ensure that sensitive data is not exposed, then it is about considering the quality, accuracy and availability of the data you gather to enable AI to improve workflows. It is also crucial that you think carefully about where the data is processed, and you are transparent about how and where it is handled. The new OneAdvanced AI software platform offers complete UK data sovereignty.

Q There are so many AI suppliers in the market – how do you evaluate who is the best partner?

A The best advice is a complete assessment of any potential suppliers' security, integration, scalability and customer support to ensure they meet your specific business requirements. Other things to consider include how data is managed and whether the supplier has experience within your sector.

Q Should we invest in a Large Language Model (LLM) when devising AI strategy?

A It can be a good idea because these are widely available, powerful and low-cost. However, as you begin the process, take time to think seriously about how you will secure your company's data. You must choose a platform that lets you safely and effectively integrate these tools into your business AI journey.

Q What's the strategic approach to deciding the extent of AI's role in your operations?

A It's scary to think that we may be the last generation to rely solely on a human workforce for routine tasks, but as a person, you get to select the future that works best for your organisation. To begin, you might focus on partially automated workflows where AI agents offer a cost-effective, low-maintenance alternative and free up your team to do higher-value work.

When you are ready, the next step is to use an agentic system where AI agents evolve from tools into collaborators and are capable of making decisions, handling tasks independently and operating with minimal human oversight.

You could then move to high autonomy, where systems manage complex workflows across departments or platforms. They will adapt to new information and adjust actions dynamically, with human input only sought when necessary.

For instance, in healthcare, document summary AI is being used by health professionals to summarise and highlight key information. The sector is also using Generative AI to draft medical documents, and AI to triage patients to save time and improve accuracy.

Eventually, you might want to have full autonomy where AI fully owns the process and operates without human intervention. This will be a future destination for many business operations, especially those requiring speed, scalability, and consistency.

Q One of the biggest risks within any secure AI strategy comes from Shadow AI – what is the advice here when educating employees?

A You must put policies in place to minimise the risks from Shadow AI. Most employees are aware that using unauthorised AI tools poses a risk to the business, whether from a cybersecurity or data governance perspective, yet many are unwilling to change how they act because of the efficiency it provides.

Commercial feature

AI: the next frontier in cyber defence

AI is rewriting the rules of cybersecurity, supercharging attackers with speed and scale, while pushing defenders to evolve or be outpaced in a high-stakes digital arms race

Artificial intelligence (AI) is transforming the rules of engagement in cybersecurity. As the technology advances at rapid speed, organisations are facing an increasingly complex digital battlefield – one where defenders and attackers alike are wielding AI as a weapon.

That's because the tools that enable speed, scale and automation for legitimate business processes are also allowing cybercriminals to launch more sophisticated, efficient and scalable attacks.

The attacker's edge: AI as a weapon

What once required human labour, for example writing phishing emails, conducting reconnaissance, probing for vulnerabilities, can now be automated and scaled with AI.

"AI allows attackers to become faster and more efficient at doing things that once required people," explains Dr Carl Windsor, CISO at Fortinet. "Tasks that were mechanical and time-consuming can now be completed autonomously. It's automated, scalable and cost-effective."

The impact is already being felt. Deepfake content, realistic phishing campaigns, malicious bots and fraudulent websites can now be generated using online services with minimal effort, dramatically lowering the barrier to entry for cybercrime. In fact, 87% of global organisations faced an AI-powered cyber attack in the past year.

The situation is made more complicated by the emergence of agentic AI, a form of AI that is goal-driven, autonomous and context-aware.

"Agentic AI has reasoning and situational awareness," says Windsor. "It can take actions based on what's happening in real time, without human input. That makes it incredibly powerful."

The dark side: AI as a business risk

But, at the same time, not all AI-related threats come from external attacks. The adoption of AI tools such as third-party large language models (LLMs) and AI cloud-based applications has created new vulnerabilities within organisations themselves. At this year's RSA Conference, one supply chain risk survey revealed that 45% of organisations using third-party LLMs experienced a security incident tied to that dependency.

"Organisations are being driven to move faster than their security teams can keep up with. The genie is out of the bottle and cannot be put back. The security team now has to deal with this new technology they don't fully control, an app, an AI model, an AI workload, previously unseen dependencies and a whole new set of security risks," says Windsor. "If you expand your attack surface so rapidly and are not fully aware of the nuances, this is where incidents can occur."

In addition, the 'shadow AI' effect – the use of unvetted AI tools by employees – may expose organisations to data leakage, model poisoning and compliance failures. Sensitive data may be fed into external models without adequate governance, creating major privacy and security risks.

"This is the difficulty," says Windsor. "AI brings with it so many benefits, it is hard to stop, but with it comes significant risk to the security of your data should it be used blindly."

Flipping the script: AI as a cyber defender

Despite these growing risks, Windsor sees an opportunity to use AI more effectively on the defence side of cybersecurity.

This approach is built on Fortinet's long-standing expertise in both cybersecurity and AI, with FortiAI representing the culmination of years of innovation and insight in defending complex digital environments. FortiAI innovations are embedded across the Fortinet security fabric platform to enhance protection against new and emerging threats, simplify and automate security and network operations and secure use of AI-enabled services and tools.

FortiAI reflects Fortinet's ongoing commitment to applying AI across its portfolio – helping organisations stay ahead of evolving threats and simplifying how they defend increasingly dynamic digital ecosystems.

"If you look at the AI usage by attackers versus defenders, we have the advantage," he says. "AI technologies can help us defend not just against AI-based attacks, but against any type of attack. AI gives us better visibility, deeper insights, faster reaction times and smarter automation."

AI can analyse vast amounts of threat data in real time, detect subtle anomalies that would go unnoticed by humans and even respond autonomously to incidents.

Fortinet's own approach integrates AI across the entire cybersecurity lifecycle, not as a bolt-on feature but as a foundational capability. Rather than isolate AI in individual tools, its FortiAI roadmap embeds intelligence across the entire cybersecurity stack.

"Our AI is part of the Fortinet security fabric platform, where its components are aware of each other. They share data. They make decisions together," says Windsor. "That's the real power – not just isolated tools, but an intelligent, coordinated system."

FortAI applies AI to the three key pillars of cybersecurity: threat intelligence, security enforcement and security operations. Each pillar has its own AI-driven focus, which together form the backbone of Fortinet's integrated security fabric.

FortiAI-Protect

FortiAI-Protect enhances detection, protection and prevention capabilities by embedding AI technologies, tools and services into both Fortinet's threat intelligence and the cybersecurity products and infrastructure enforcing security. The goal is faster, more accurate identification and mitigation of known and unknown threats, whether conventional or AI-driven.

"AI helps us recognise threats in real time, understand them more deeply and react accordingly," says Windsor. "We can also detect AI-specific risks, like synthetic content or model misuse."

FortiAI-Protect gives security teams improved insight into the threat landscape and speeds up incident detection and response across Fortinet solutions, on-premise or in the cloud.

FortiAI-Assist

FortiAI-Assist combines AI-driven analytics and automation for security and network operations – critical areas plagued by talent shortages and increasing complexity.

"How do we make the Security Operations Centre (SOC) and Network Operation Center (NOC) more efficient and effective? By using AI to analyse incidents, automate configurations, actively hunt for threats, analyse the network, applications and data and autonomously take actions to protect, mitigate and optimise availability and performance," says Windsor.

The platform uses generative and agentic AI to provide deep and correlated analysis, suggest and automate network and security configurations and adjust systems dynamically in real time. Fortinet tools can now create and implement network changes in minutes that used to take hours or days.

FortiAI-SecureAI

Finally, FortiAI-SecureAI is Fortinet's answer to the growing risks posed by enterprises' use of AI tools. It secures AI models, workloads and their underlying infrastructure – preventing data leakage from LLMs and ensuring data integrity. It addresses issues such as prompt manipulation, shadow AI, data leakage, supply chain risks, data and model poisoning and zero-trust access.

"AI tools used by your employees – internally or externally – become part of your attack surface," says Windsor. "FortiAI-SecureAI tools and capabilities ensure that this usage is vetted, monitored and protected."

AI as a force multiplier

The growing role of AI in cybersecurity is inevitable. But Windsor argues that the outcome depends on how organisations harness it.

"Any technology can be used for good or bad," says Windsor. "And, like any powerful technology, AI can be dangerous if misused. But when smartly used throughout the cybersecurity stack by defenders, it's a force multiplier."

Ultimately, as AI accelerates the pace and scale of cyber threats, the stakes for organisations have never been higher. To stay ahead on this fast-evolving battlefield, businesses must treat AI not just as a defensive add-on, but as a strategic capability woven into the very fabric of their cybersecurity approach.

> “Organisations are facing an increasingly complex digital battlefield

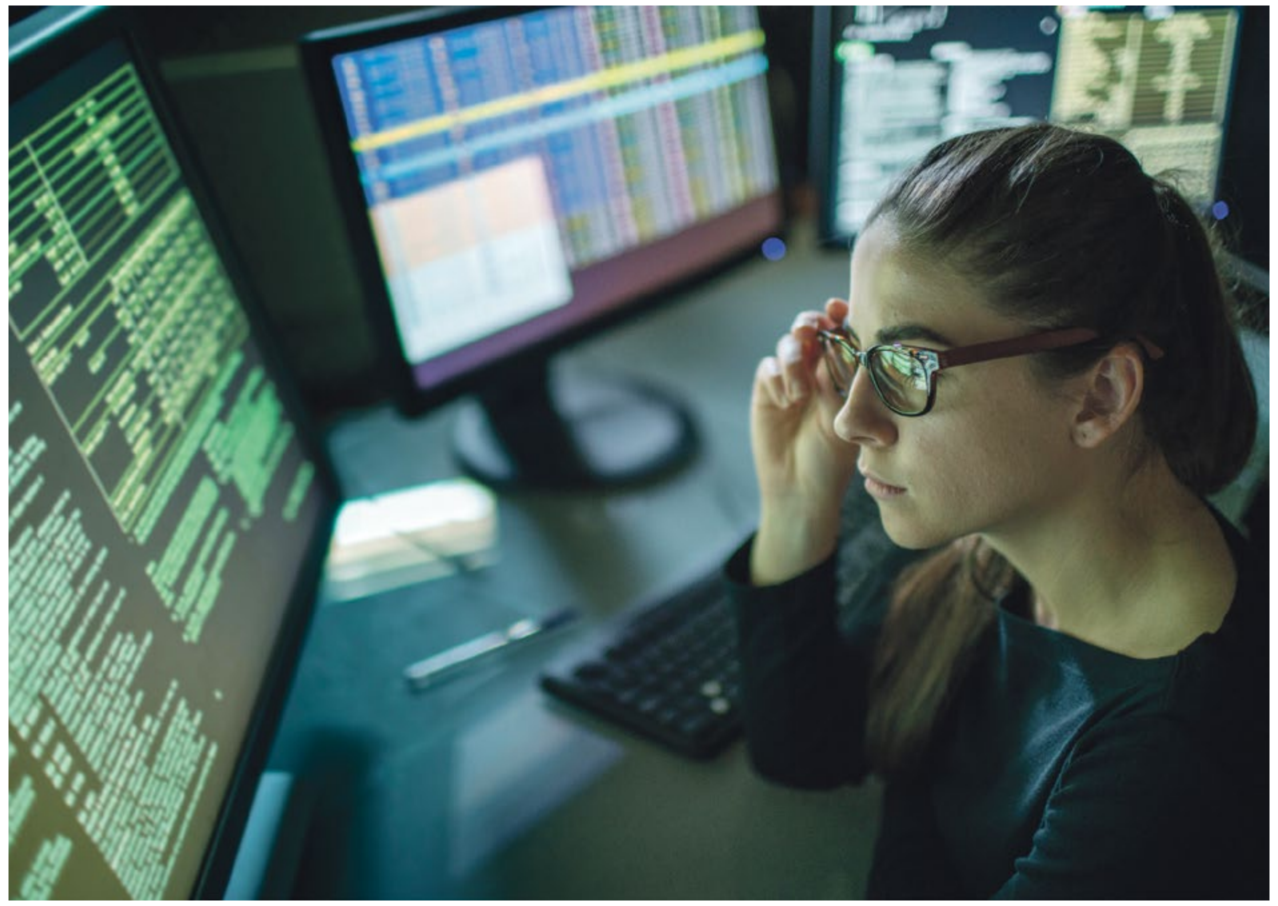

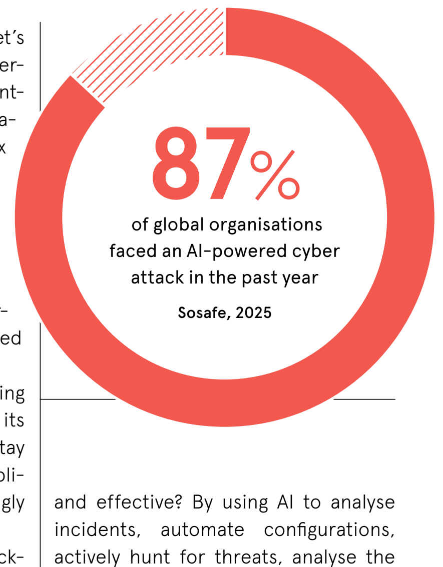

87%

of global organisations faced an AI-powered cyber attack in the past year

Sosafe, 2025

Listen to *Brass Tacks* – a Fortinet podcast that gets straight to the point on cyber risk, made for executives at fortinet.com/resources/podcast

AI AVATARS

Chief execs deploy digital doubles to take their meetings

Some executives are using AI avatars to deliver important updates. Could these experiments serve as a proof-of-concept for bot-led company all-hands?

Sam Forsdick

One of the benefits of AI, we're often told, is that it will relieve us of workplace drudgery by automating menial tasks. For senior leaders, one particularly tedious task may be ripe for automation: meetings.

According to a Harvard Business School study of 27 CEOs and 60,000 hours of work, chief executives spend 72% of their working days stuck in meetings. Not all of these meetings are useless, of course. But many C-suiters would happily dedicate more hours to shaping strategy, nurturing culture or coming up with new product ideas, if they could only reclaim some of their time.

Executives in the tech sector have trialled the use of AI avatars as stand-ins in the meeting room. The CEO and co-founder of Klarna, Sebastian Siemiatkowski, sent an AI double to speak at a recent financial update. Sam Liang, the CEO of Otter.ai, a transcription software company, has developed a 'Sam-bot', which will eventually take his place at company meetings. And Eric Yuan, the CEO of Zoom, used an AI avatar to deliver the initial comments in the company's Q1 earnings webinar in May. Yuan, who plans to use his digital double in most meetings by next year, envisions a future where the avatar will even make decisions on his behalf.

Steve Rafferty, head of APAC and EMEA at Zoom, has also experimented with using an AI avatar in meetings. While the technology is not yet ready to interact with other people or answer questions, it has helped him deliver more personal messages to employees.

"My team stretches from the Arctic Circle all the way to Antarctica," he says. "And there's roughly 60 different languages spoken across those regions." Hosting meetings therefore can be challenging. In April, Rafferty used his AI avatar to introduce a quarterly meeting in fluent French. "It means I can be in multiple places at once, speaking different languages," he says.

Rafferty describes the tool as "another string to the bow for business communication" and claims his teams have bought into receiving messages from his AI double.

To help the technology progress, Zoom is training its executives' AI companions on their communication and decision-making style. "I always have my AI companion turned on," Rafferty says. "It's across my messages, video, phone and chat. It's across everything." The next challenge, he adds, is trusting the AI to make decisions.

Dan Thompson is the CEO of Sensay, a startup that develops AI replicas of human employees. These avatars are trained on the individual's calls, messages, emails and documents. By learning how the individuals speak and write, the bot can surmise their likely response to some questions. Thompson uses his own AI replica to draft emails and messages, which he estimates saves him from hours of administrative work.

He explains: "I was sorting my visa at the embassy this morning, which took two hours. While I was there, all my morning emails had already been pre-drafted and all I had to do was read them and decide whether to hit send or edit."

While there are some obvious advantages to using AI in this way, doing so also presents unique risks. AI systems are prone to hallucinations and it is not uncommon for them to present misleading or inaccurate information as facts. Such errors could be detrimental coming from a company CEO.

Security is also a concern. Last year, fraudsters used deepfake technology to imitate the voice of Mark Read, the CEO of WPP, a multinational advertising company, and attempted to solicit money and personal information from the leader of a WPP agency. Employees could become more susceptible to such scams if they come to expect important messages to be delivered by AI avatars.

Rafferty says it's essential that businesses develop a strategy for this technology, detailing what kinds of communication people should or should not expect from an AI companion. "You can't just go blindly into this," he adds. "A lot of people are selling AI as the answer to everything, where it actually could be the problem. You need proper governance, structure and processes in place so people can trust it."

To train these executive AI doubles, firms must give large language models access to sensitive company information. Rafferty, for example, has allowed his AI companion to access his work phone and recordings of his meetings.

But information can easily be leaked if robust security protocols are not established. For instance, last year a Microsoft employee warned that staff could use the Copilot AI platform to access their colleagues' HR documents or read executives' emails.

According to Rafferty, these security risks result from human error rather than inherent vulnerabilities in the technology. "That's down to the way the platform is set up," he says. "It can be told to share only data that's relevant to people at a certain level in the business or make it available to everyone. But, if the planning is poor, you can expose yourself."

AI bots may be fine for handling low-stakes tasks, but they are unlikely to replace other elements of the CEO role no matter how advanced they become. "The job of a CEO isn't just about outputs – it's about meaning. AI can't embody organisational purpose, grapple with ethical dilemmas or inspire people through shared struggle." So says Dr Alexandra Dobra-Kiel, director of innovation and strategy at Behave, a consultancy.

And certain types of business decisions require a level of nuance that bots cannot provide. "AI may be able to analyse trends, but it can't navigate the ambiguities of long-term vision, weigh trade-offs appropriately or mimic the instincts that define high-stakes leadership," Dobra-Kiel adds.

So while Yuan and other executives may offload some of their meetings to AI assistants, it is unlikely that they will delegate their decision-making to digital doubles any time soon. Emphasising the importance of human leadership, Dobra-Kiel says: "People follow leaders they believe in, not just voices that sound like them." ●

> “People follow leaders they believe in, not just voices that sound like them

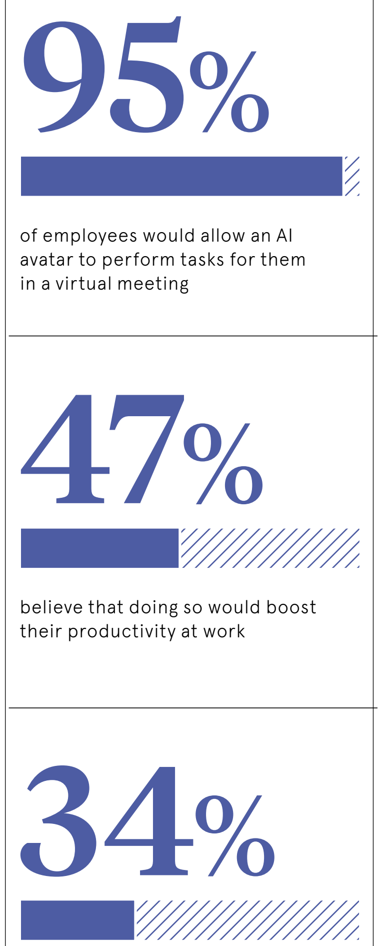

95%

of employees would allow an AI avatar to perform tasks for them in a virtual meeting

47%

believe that doing so would boost their productivity at work

34%

of executives would allow an AI avatar to make decisions on their behalf

Travelperk, 2024

THE STATE OF AI ADOPTION

Organisations have been hastily adopting AI tools for several years, but most have struggled to integrate the technology into their operations effectively. That is beginning to change. Best practices for AI integration are developing rapidly and more firms than ever before are achieving significant benefits with AI tools. However, questions about the wider impacts of AI adoption must be addressed.

AI ACROSS THE ORGANISATION

Share of organisations worldwide using GenAI in the following business functions

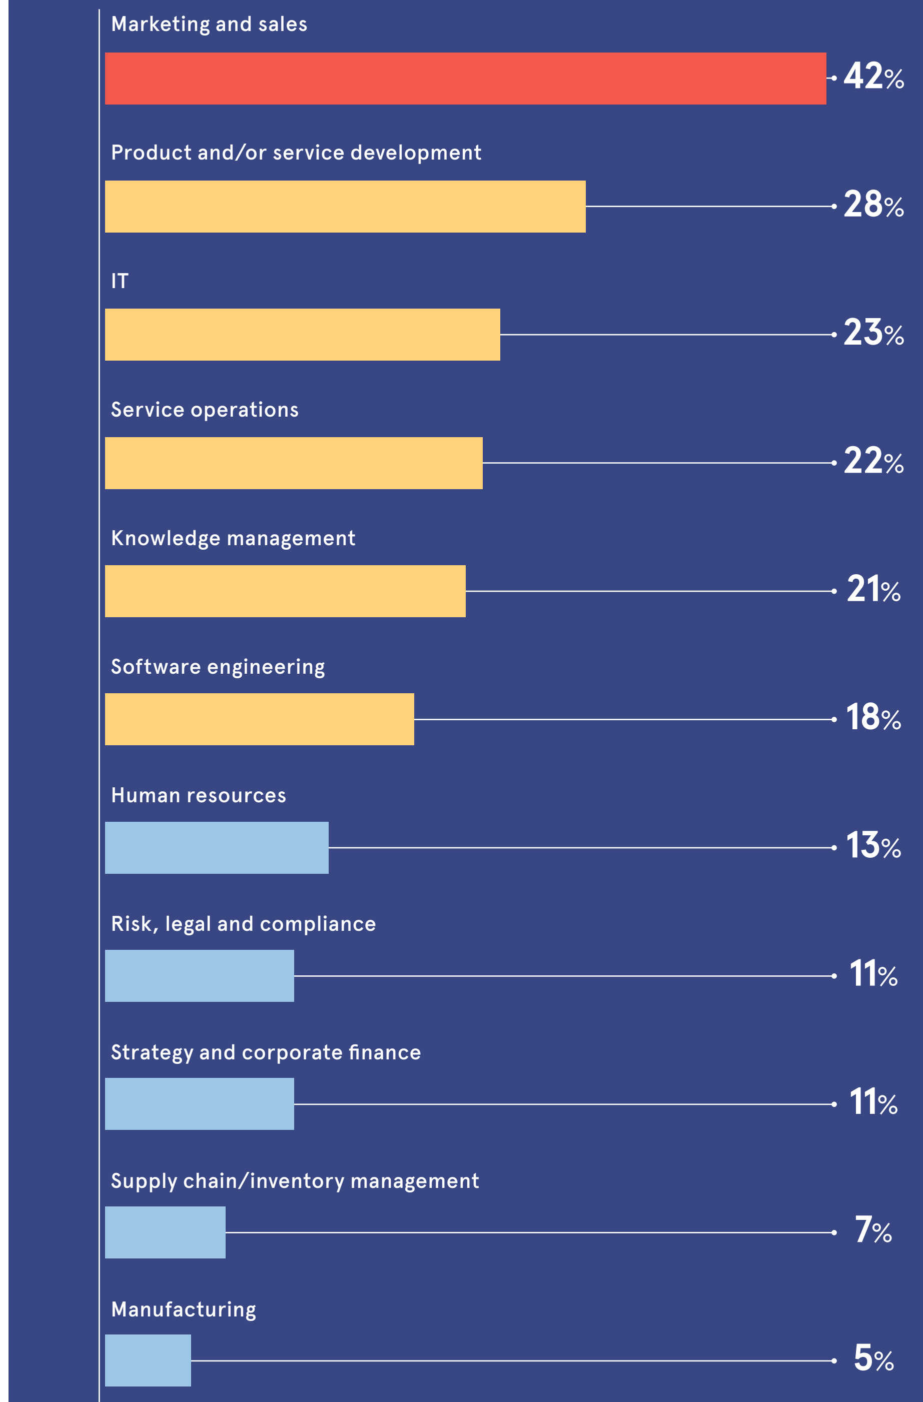

FIRMS HAVE ACCELERATED THEIR USE OF AI SINCE 2023

Share of organisations using GenAI in functions across the business

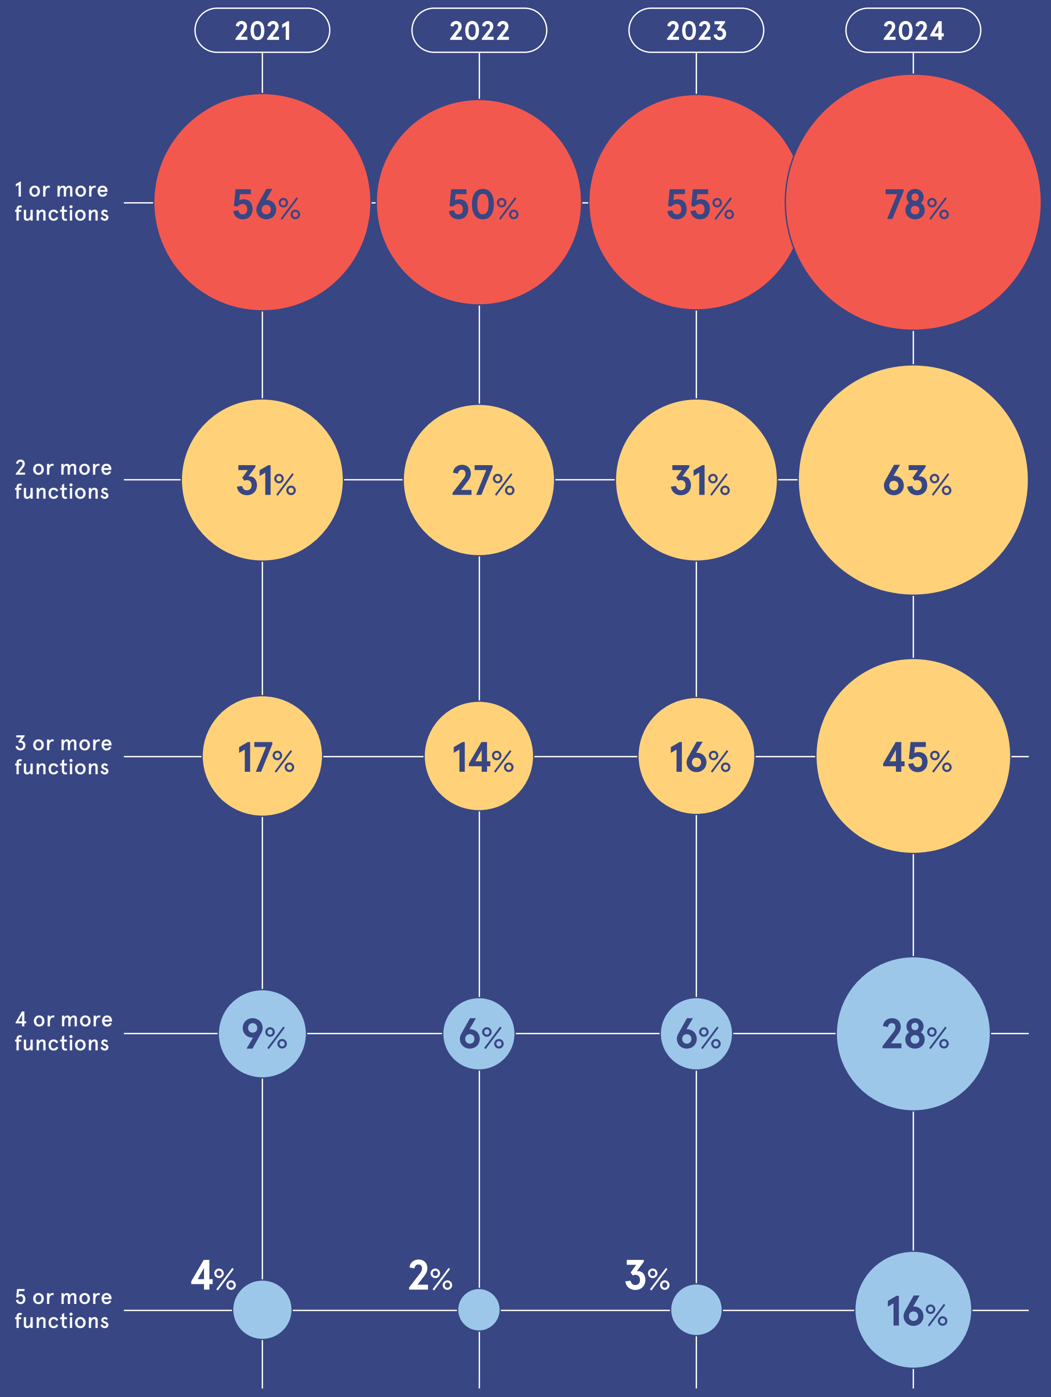

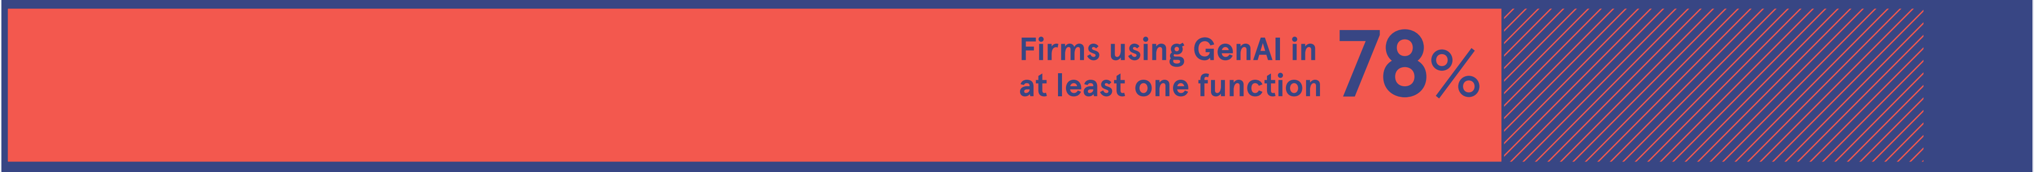

BEST PRACTICES FOR ENTERPRISE AI ADOPTION

Share of organisations adopting the following best practices, by size

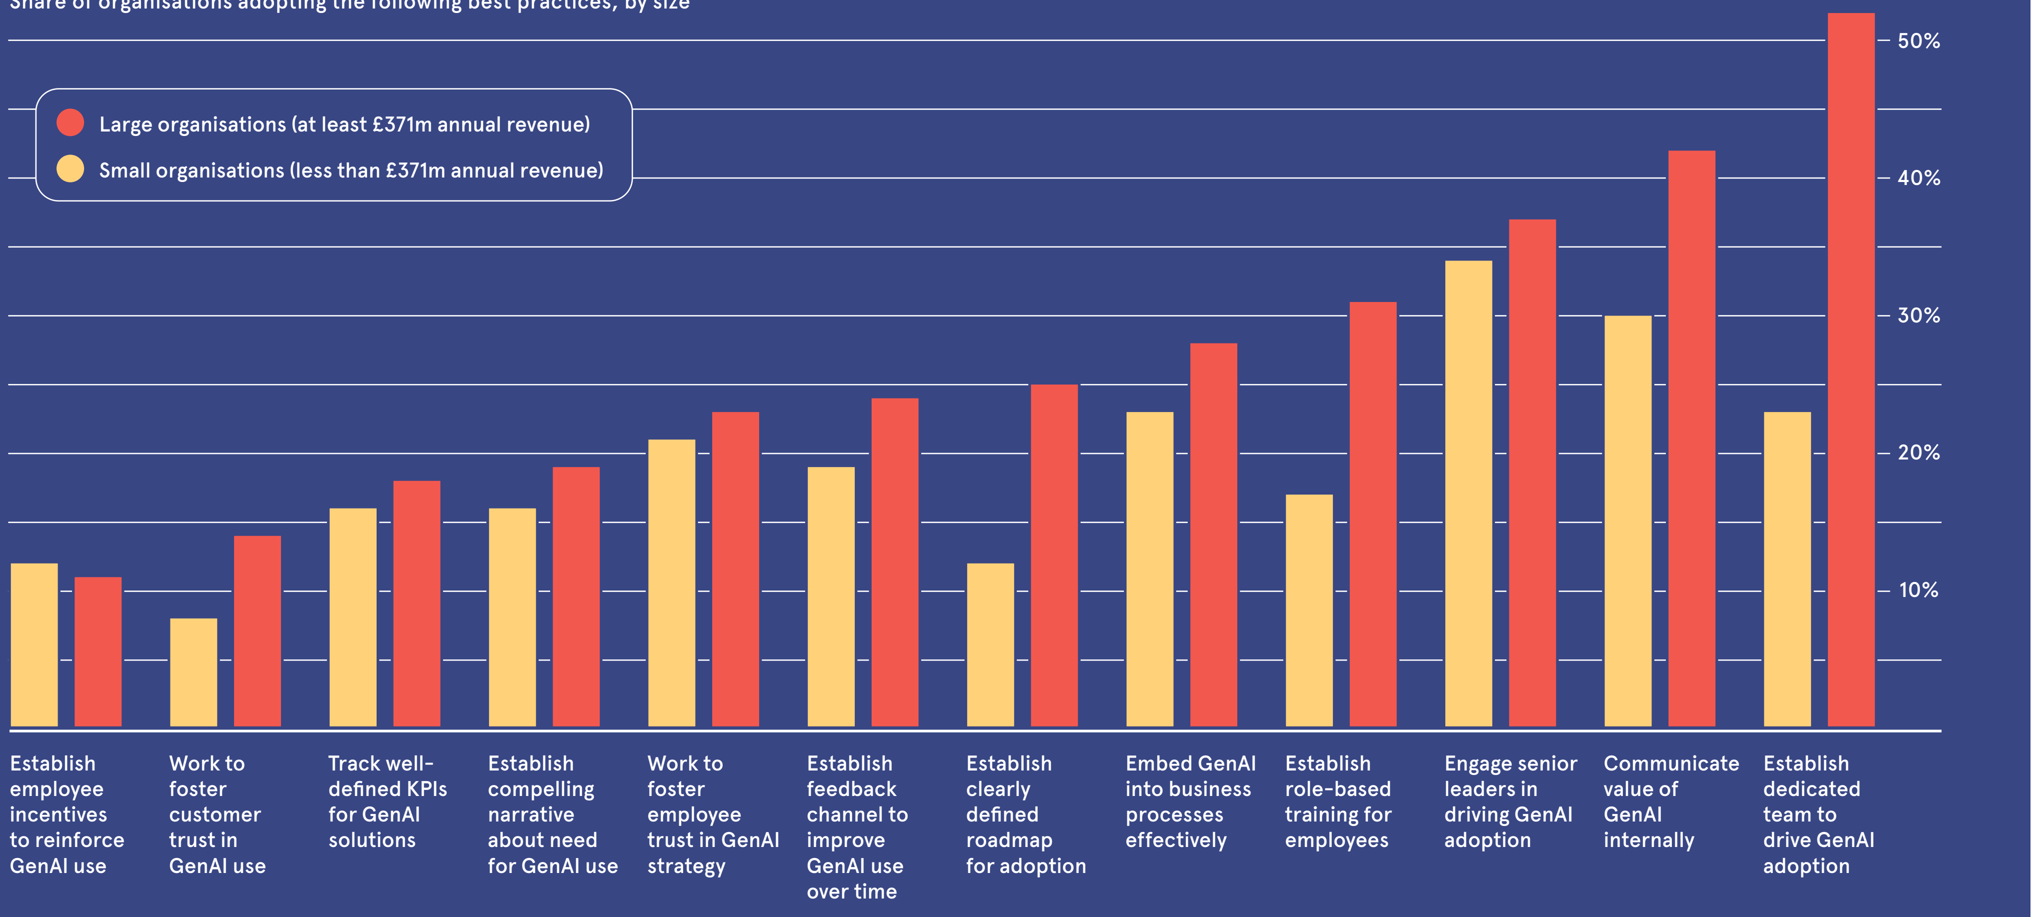

MOST FIRMS HAVE DEVELOPED A CENTRAL HUB RESPONSIBLE FOR AT LEAST SOME AI DECISION-MAKING

Share of organisations structuring AI responsibilities in the following ways

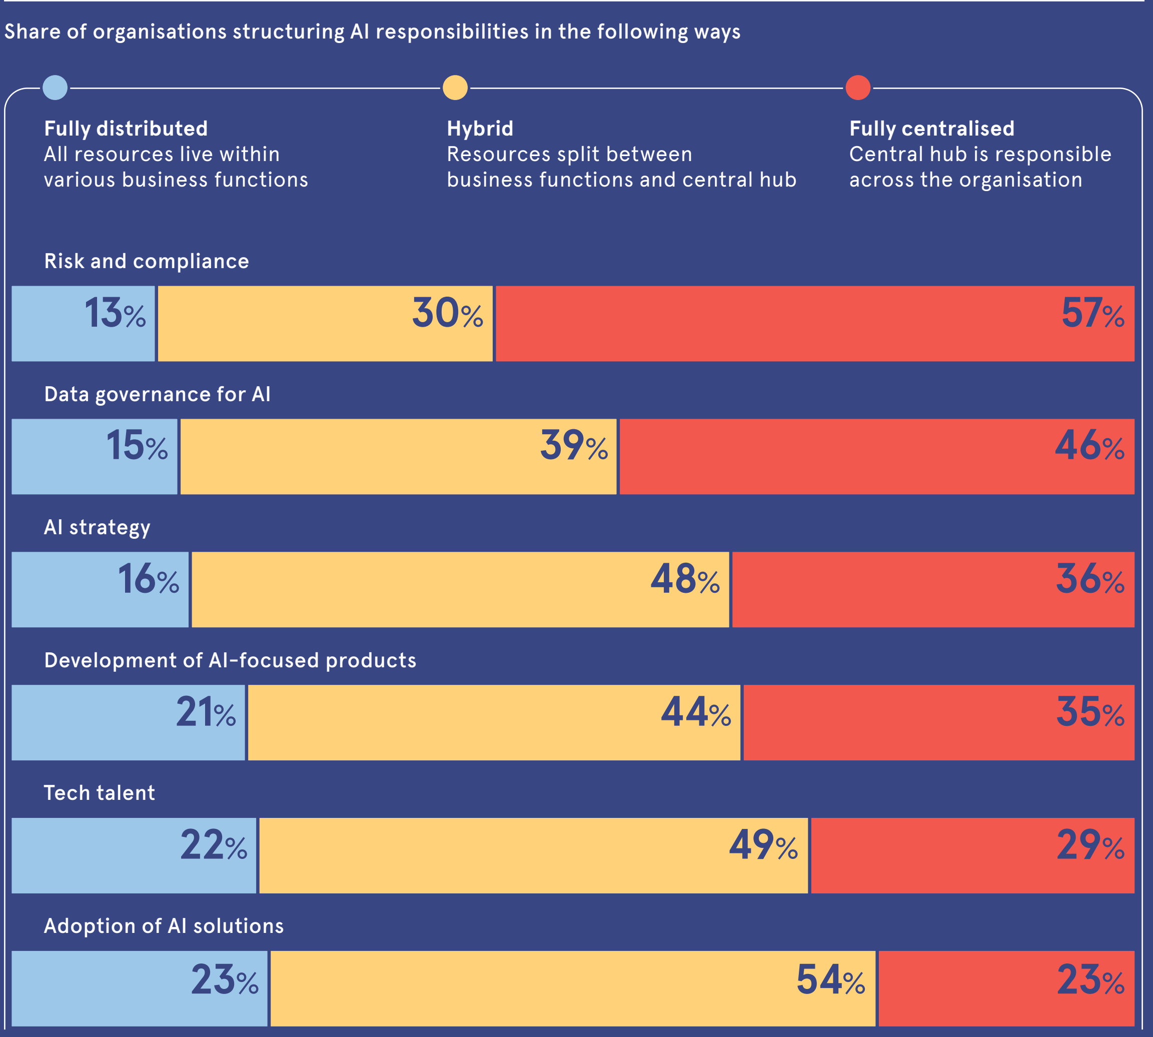

RESKILLING OR REDUNDANCY

Share of organisations planning to reskill portions of the workforce as a result of AI adoption by 2027

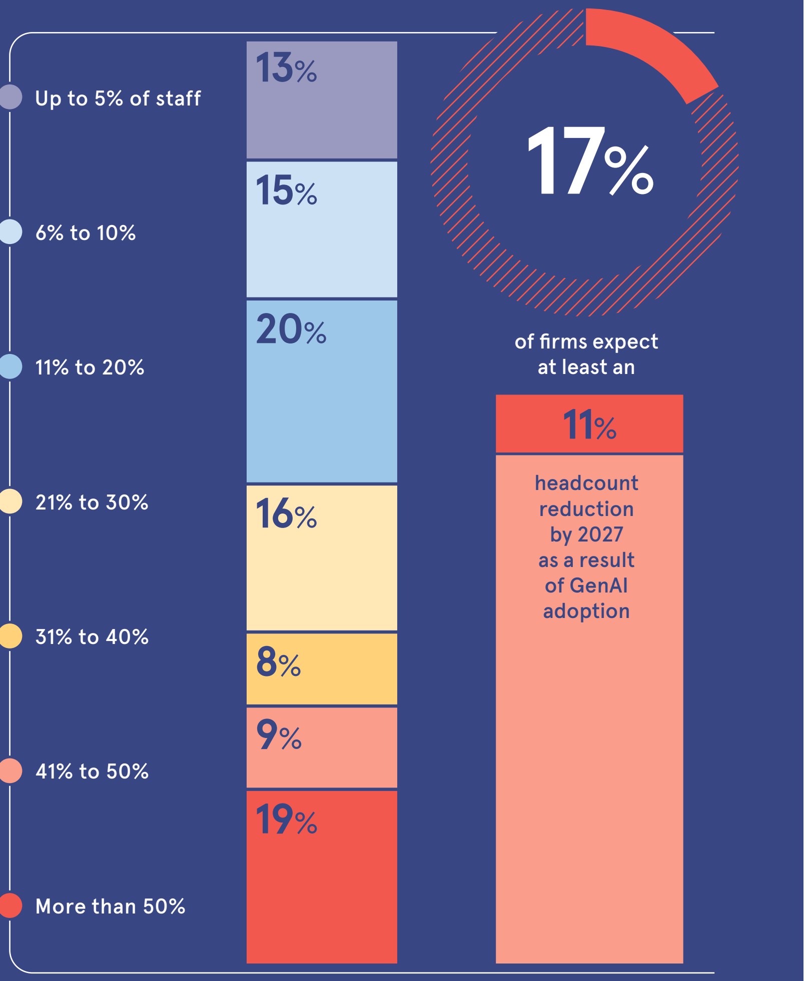

McKinsey, 2025

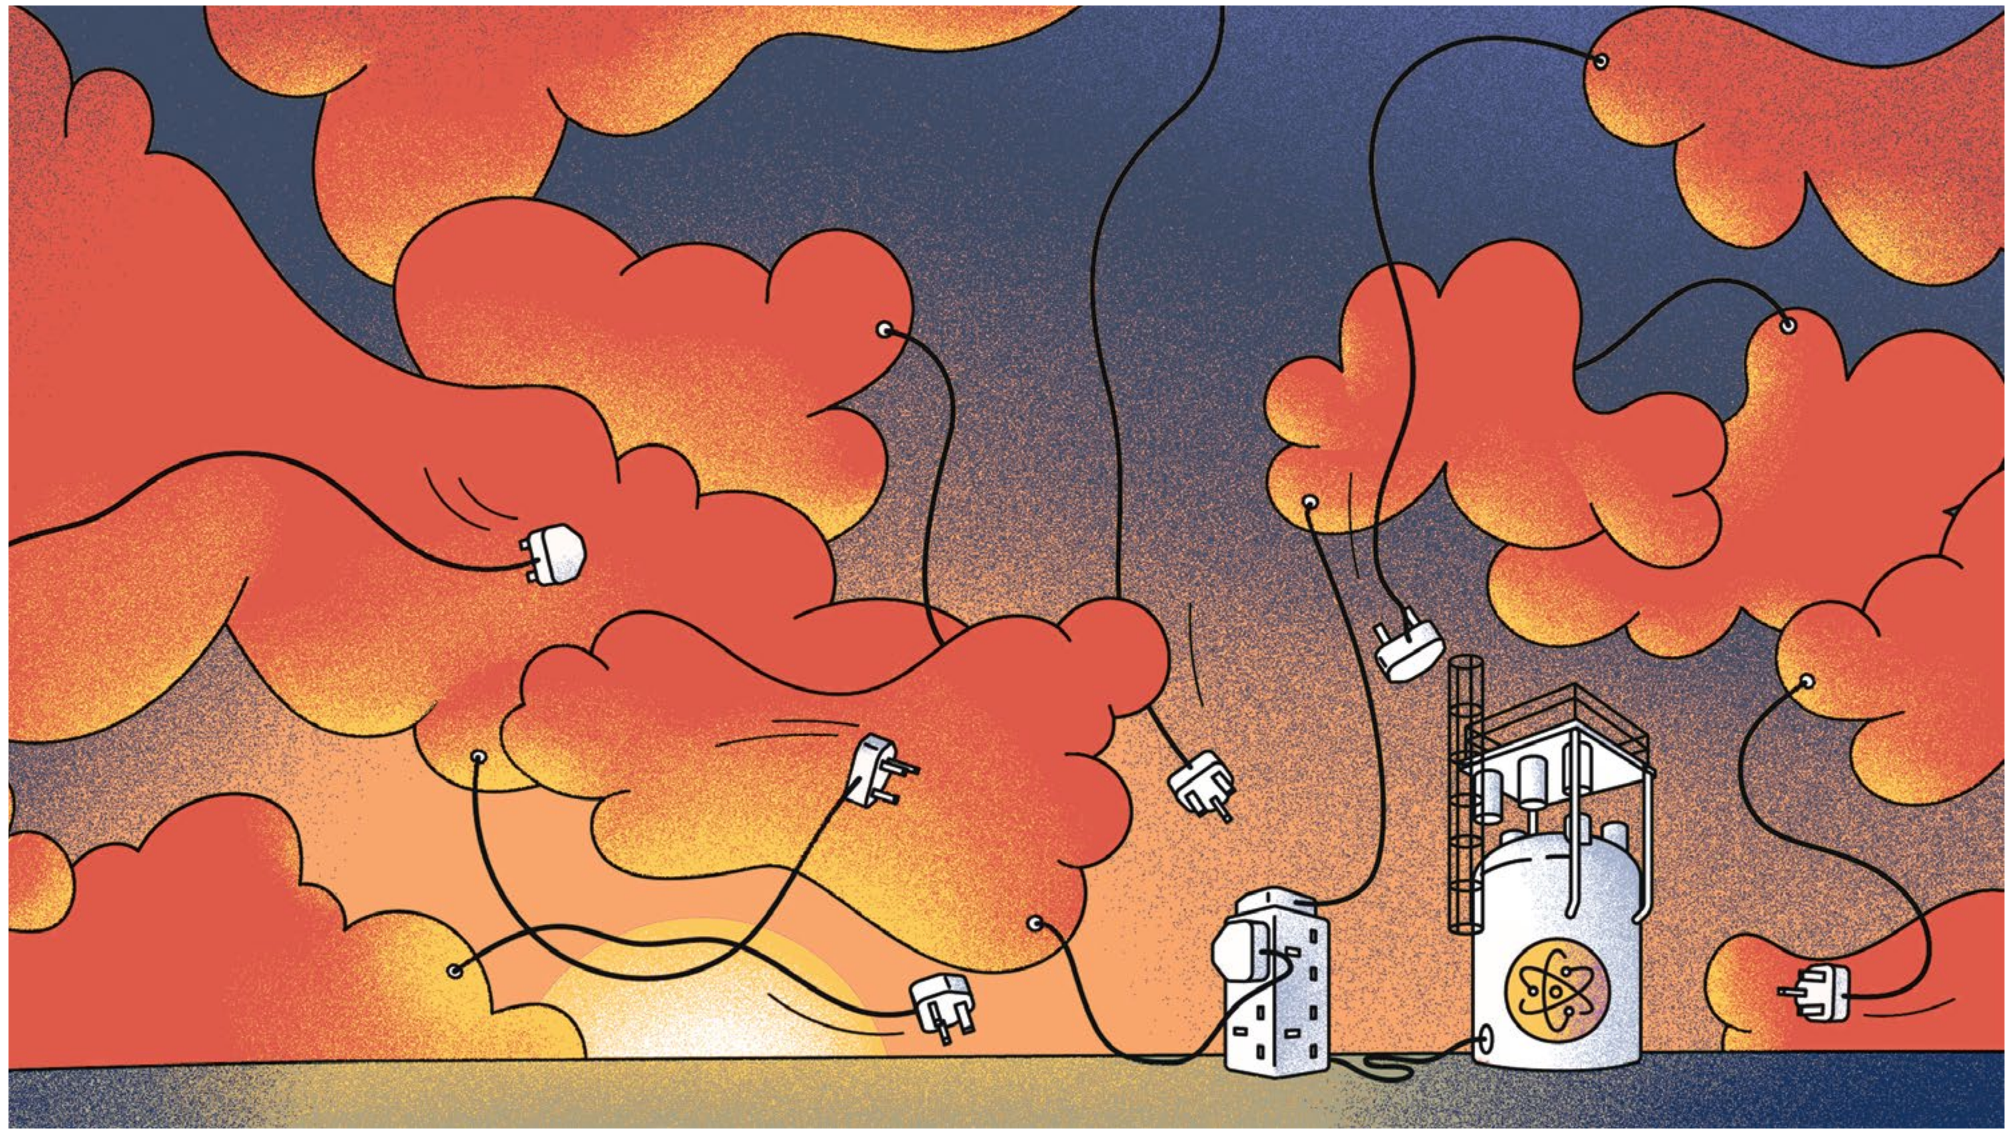

DATA CENTRES

Challenges abound for nuclear-powered data centres

Keen to ignite an AI revolution in the UK, the government has relaxed planning laws for a novel type of nuclear reactor, which could be used to power data centres

Tamlin Magee

When the Labour Party came to power in 2024, Keir Starmer immediately opened the gates for a data centre construction boom. The move wasn't just a way to attract investment in a shaky economy – it was part of a larger strategy to make the UK a leader in AI technology.

AI systems are fuelled by data, so as these systems proliferate, more data centres will be needed to support them. But the challenge then becomes fuelling those data centres, which require significant energy resources to operate.

In the UK, data centres already use 2% of total power supplies. By the end of 2025, data centres worldwide could guzzle as much as 23 gigawatts of power – or twice the amount consumed by the Netherlands. Lacking suitable infrastructure to support their soaring resource demands, data centres are competing with population centres for power supplies.

Elevated energy costs are also threatening future data centre operations. Wholesale electricity in the UK is already expensive compared with peer countries in the EU. And energy prices are still higher than in pre-Covid years for most developed economies.

The world's largest tech companies are therefore seeking alternative energy sources to power their data centres. With clean-tech solutions unable to meet the power demands of energy-hungry industries, nuclear power is drawing interest from the public and private sectors. And key to making nuclear energy more accessible for commercial use is a nascent technology that may finally be ready for production: small modular reactors (SMRs).

SMRs are designed to generate less than a third of the power created by traditional nuclear power stations. But they take up far less space, which is partly why the industry is excited about their potential.

"Data centres are probably the best market signal for these technologies," according to Ross Peel, research fellow for nuclear safety and risk at King's College London.

Westminster has committed to exploring the use of nuclear power to fuel data centres. And SMRs will be key to enabling growth in the country's AI industry, according to a spokesperson for the Department of Energy Security and Net Zero.

Great British Energy Nuclear recently selected Rolls-Royce to build the UK's first SMR. The reactor won't be operational until the 2030s, but Ed Miliband, the energy secretary, said the project marks an end to the "no-nuclear status quo".

Nuclear projects in the UK are regulated by several departments. The Environment Agency, the Department for Environment, Food and Rural Affairs and the Planning Inspectorate are responsible for the planning and regulatory justification of new sites. The Office for Nuclear Regulation (ONR) assesses the safety and technical standards of nuclear providers.

Although the ONR is not directly responsible for setting policy, Jane Bowie, its director of new reactors, says the organisation has been working for years to streamline the regulatory process.

Now that large sites such as Hinkley Point C have been approved, the regulators have turned their focus to SMRs. The UK has not yet approved any SMR designs. But applications may start pouring in if approvals processes are streamlined.

According to Chris Gadomski, chief nuclear analyst at BloombergNEF, the industry is waiting at the starting gates. "All these providers are lined up and boom – when the rate goes down, if it's approved, you will have an onslaught of advanced-reactor companies moving through the licensing process expeditiously."

It's for good reason that regulatory approvals for nuclear projects take so long. Faults or errors in the design or operation of nuclear facilities can be catastrophic. Any new nuclear facility must pass a meticulous design-approval process before it is built. Once it is operational, the site must comply with strict maintenance and security rules.

Crucially, any waste must be disposed of appropriately. It cannot be discarded or destroyed; it must be stored, usually underground, until it is no longer harmful. Some types of waste, such as plutonium-239, have a half-life of 24,000 years.

Waste storage will become more urgent if SMRs gain regulatory approval. Not only would more firms rely on nuclear power to fuel their operations, but SMRs might also produce relatively more waste than conventional reactors, according to research by Stanford University and the University of British Columbia.

> "Realistically, SMRs won't be available fast enough to solve the near-term power crunch

"Who will own the SMRs used to power data centres?" asks Lindsay Krall, a geochemist and lead author of the research. "And will those owners be prepared to manage the nuclear waste they will produce?"

Authorities must consider these questions during the technology-selection process, she adds, long before deploying any SMRs.

Peel says the UK is capable of storing nuclear waste, but finding acceptable sites to host it has been challenging. "The main issues are political," he says. "Until you get a site, there's not much we can do. We can research materials, but without a site it's all just academic."

> "Until you get a site, there's not much we can do. We can research materials, but without a site it's all just academic

The thought of streamlining nuclear regulation may raise some eyebrows. For obvious reasons, the industry devotes considerable resources to safety and security. While these are still priorities for nuclear authorities, pressure is mounting to get developments off the ground.

A culture clash between tech providers and the nuclear industry is unfolding. Peel says regulators want to be involved in the design of new nuclear technologies, and they want operators to develop reactors slowly. But developers, and even some policymakers, are pushing for more flexibility.

But the fail-fast model preferred by some startup providers simply won't work in the nuclear industry, says Peel. "You can't have your products just fail."

Small SMR firms may also struggle to meet the industry's strict security requirements. Some providers, says Peel, envision their facilities as largely autonomous sites, which are connected to the internet and can operate with minimal human oversight. But nuclear sites need physical and virtual protections.

Conventional reactors used for non-military purposes are usually policed by armed security forces maintaining tight perimeter security. Providers hope to avoid hiring such forces by clustering their sites around existing plants. But part of the appeal of SMRs is that they can be widely distributed.

Cyber attacks too threaten nuclear sites. Because power infrastructure has proved a tempting target for cybercriminals, nuclear regulators are crafting new digital-security standards for SMRs.

Providers are attempting to mitigate cyber risks in the design of their reactors by, for instance, developing materials that prevent reactors from reaching meltdown temperatures.

But Peel is unconvinced that such innovations will solve the security problem. "Every nuclear disaster – Fukushima or Chernobyl or Three Mile Island – happened in a way that nobody foresaw," he says. "It's not the accident you've planned for that gets you, it's the one that no one thought about."

Data centre operators are also pushing for faster regulatory approvals. But even if their lobbying efforts are successful, new nuclear reactors will not arrive in time to satisfy near-term capacity needs.

"Any notable nuclear contribution to the grid is going to be after 2030. But hyperscalers needed the electricity yesterday," Gadomski says.

According to Patrick Smith, field chief technology offer at Pure Storage, a hybrid-cloud integrator, energy-efficient LLMs won't solve the problem either. "As systems become more efficient, people just do more with them," he explains.

He continues: "Realistically, SMRs won't be available fast enough to solve the near-term power crunch, particularly in markets such as the UK, where AI growth is outpacing infrastructure readiness. Firms therefore will continue to rely on renewables and fossil fuels for the foreseeable future."

So while the promise of sustainable energy might motivate government and industry to develop and enable new nuclear technologies, it will be many years until these reactors are widely used. For now, SMR developers are facing regulatory hurdles, logistical problems and security questions, not to mention a general resistance to nuclear energy among the public. The road to a nuclear-powered future, it seems, will be long. ●

AI SYSTEMS ARE INCREASINGLY ENERGY EFFICIENT

Energy intensity of AI computer chips; index score, log scale

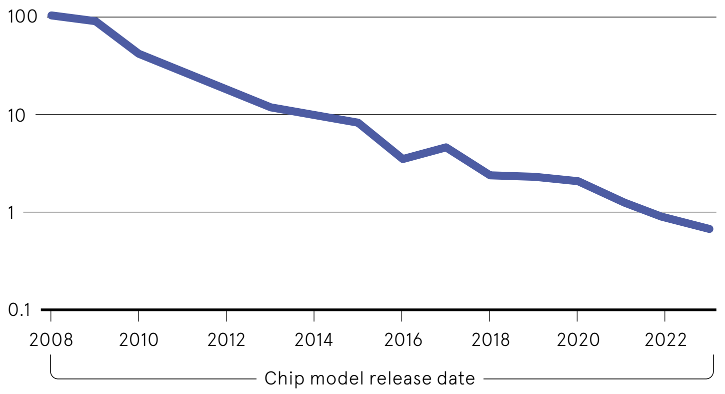

International Energy Agency, 2024

Commercial feature

AI in financial services: from internal gains to external value

As AI agents increasingly handle customer decisions, from shopping to investments, firms must compete for AI's attention rather than human loyalty, according to industry experts

What happens when the cost of intelligence sinks to zero? This question, posed at a recent industry roundtable hosted by Cognizant and Raconteur, cuts to the heart of the financial services industry's transformation challenge.

When artificial intelligence can write reports, conduct research and generate ideas instantaneously, traditional ways of working – including customer relationships – face fundamental disruption. The winners will be those who move the fastest to transform AI from a back-office tool to a front-line competitive advantage. Those quick enough to harness AI's shift from internal efficiency to external value creation will find themselves competing in an entirely new game.

The race is already underway. By 2030, AI-friendly consumers are expected to drive 55% of all purchases, according to recent Cognizant research, equivalent to £690bn in UK consumption alone, which will fundamentally change how the financial services sector must operate.

For financial institutions, this will mean competing for the attention of customers who use AI for interaction. When AI agents evaluate mortgage rates or process loan applications at machine speed, traditional customer touchpoints become obsolete.

The solution lies in what David Fearne, director of generative AI at Cognizant, calls the "agentic enterprise" – moving beyond isolated AI tools to interconnected systems that mirror human organisational structures, with enterprise intelligence coordinating task-based AI across business units. This breaks down the walls between individual AI applications, creating a unified intelligence layer that can pursue opportunities too marginal for traditional approaches to handle.

Notably, UK financial firms are responding with investment: 16% of technology budgets now go to AI, up from 11.7% last year, FinTech Magazine reports. Yet Cognizant research has found that only 30% of BFSI organisations have rolled out cross-enterprise GenAI use cases.

Speed becomes survival

The transformation's velocity defies traditional planning cycles. At Pismo, CEO Vishal Dalal describes the pace as "whiplash-inducing". "A few months ago, we were doing the basics well with Copilot," he explains. "Then suddenly one customer said: 'Show me how you can deliver value with AI, and we were off to the races."

This stinging comment spurred a scattering of AI experimentation that is already bearing fruit. Out of Pismo's last 10 client proposals, five have demanded AI capabilities – a significant increase in just a month. Core banking transformations that historically required years can now compress processes from months to minutes. Dalal's team, in a recent demonstration, were able to achieve deposit configurations in minutes, when it previously took several weeks to complete – all thanks to automation.

Andrew Bateman, executive vice-president of lending at Finastra, sees similar acceleration. The company's survey of 1,100 senior executives, across 11 countries, revealed 61% of financial institutions deployed or improved their AI capabilities last year, nearly doubling from 37% in 2023. "If you don't adopt now, you're still going to be taking six to nine months to get an idea out, whereas your peers will be doing it in weeks," he warns.

Andrew Shannon, global head of IT infrastructure at TP ICAP, observes the evolution from chatbots to collaborative agents. "We've started to evolve to agents that can act on tasks, be part of teams and work together," he explains. "There are huge opportunities to re-engineer how organisations work."

> Clients who began AI initiatives in 2023 are leaps and bounds ahead of their competitors

Building enthusiasm over resistance

Cultural transformation proves as crucial as technology. TP ICAP deliberately started with AI productivity applications in lower-risk environments before expanding into revenue generation. Now, by encouraging an "AI culture" and recently establishing an "AI and Innovation Lab", the company has sparked widespread interest in exploring possibilities. "I'm inundated by enthusiastic people from across the business who want to understand opportunities," says Shannon.

Elsewhere, Finastra hosts "GenAI expos" featuring "prompt-a-thons" – an evolution of hackathons that builds repositories of effective natural language instructions. "Getting people to think about how to have that interaction is sometimes the hardest thing," explains Bateman. "You want natural language rather than programmatic approaches."

However, resistance persists. Fearne encounters "unhealthy levels of scepticism" across financial services, primarily driven by fear, from clients and their employees. His solution? "Write internal charters defining what AI will and won't do."

Indeed, governance and regulatory concerns compound the challenges, although attitudes are likely to shift rapidly as competitive pressures intensify. Blue Prism research, published earlier this year, shows that 76% of financial firms plan to implement agentic AI within the next 12 months, while 40% of consumers would trust AI to help them learn about financial services, according to Mintel. Both percentages look set to rise quickly.

The talent landscape transforms, too. Shannon notes that certain generations will soon only accept jobs that utilise AI tools. Companies that restrict usage risk excluding top talent, particularly as all graduate interns now use ChatGPT and similar systems.

Rethinking skills and systems

What is increasingly clear is that many employees require immediate reskilling to adapt to AI-driven changes. As a result, traditional technical skills are giving way to human capabilities, including emotional intelligence, strategic thinking and effective communication.

Fearne's team hired their first psychology graduate as an "AI psychologist" to understand how sophisticated models respond to inputs. "These models are so complicated that even big labs can't understand what's going on," he explains. "We study them by cause and effect, the same way we study humans."

Financial services recruitment shifts accordingly, prioritising communication skills over traditional computer science backgrounds. "Natural language is the new programming language," notes Fearne. "Everyone's going to become an AI boss of some description."

> If you don't adopt now, you're still going to be taking six to nine months to get an idea out, whereas your peers will be doing it in weeks

Technical infrastructure and partnerships, as well as data quality and management, remain vital for financial services firms to stay ahead with AI innovation and, more crucially, remain relevant. Fearne says, for example, that Microsoft's Azure AI Foundry provides the platform for this enterprise intelligence layer, integrating across multiple systems rather than being limited to single applications like traditional vendor AI. "This enables the strategic coordination that makes agentic enterprises possible," he adds.

Early movers are enjoying substantial advantages. "The clients who begin AI initiatives in 2023 are leaps and bounds ahead of their competitors," says Fearne. "They took the plunge, learned repeatedly and now that understanding is paying ridiculous dividends."

Leadership advice for financial services leaders seeking to improve AI capabilities and capture external value centres on three principles. First, act immediately. "Don't wait," urges Dalal. Fearne adds: "Be careful not to let perfection be the enemy of good. Businesses operate successfully with imperfect humans."

Second, maintain what Dalal calls "curiosity on steroids". He advises that leaders adopt this mindset when future-proofing, focusing on the changes needed today to remain competitive as the industry rapidly evolves. Third, embrace the learning curve that early movers now enjoy as a competitive advantage.

"We're going to see the individual use cases that people have been working on start to become increasingly interconnected with AIs talking to other AIs and an order starting to emerge as the agentic enterprise fully emerges," adds Fearne.

Ultimately, financial institutions are at an inflexion point. Success requires combining technological capability with human insight, creating AI-enabled cultures while maintaining trust. In turn, institutions must recognise that those who master AI will compete with each other on a higher level in tomorrow's market, while those who hesitate will find themselves scrambling for relevance in today's market.

For more information please visit cognizant.com

cognizant

Microsoft

WORKFORCE

Will AI replace roles in the UK civil service?

The UK government plans to reduce costs and headcount in the civil service by using AI to automate tasks. But a transformation project targeting role replacement could be doomed for the outset

Tamlin Magee

Keir Starmer and Peter Kyle, the science secretary, have outlined the government's intention to transform the UK civil service with AI. They hope that automation will not only make public services more efficient, but also enable government agencies to reduce headcount.

Starmer declared in March that he is "determined to seize" the "golden opportunity of artificial intelligence" to deliver cost savings for the UK. And in an interview with Sky News, Kyle said explicitly that digitalisation and AI would likely lead to job cuts in the civil service.

According to Kyle, too many government employees are spending their time on administrative tasks that can be automated by AI, such as answering phone calls or sorting physical documents.

These statements are unusually candid. Messaging around AI typically focuses on the technology's power to aid human workers rather than replace them.

The Public and Commercial Services Union, which represents workers in the UK government and public bodies, has acknowledged that technology can improve public services but maintains that AI "cannot be used as a blunt instrument to cut jobs".

But many industry leaders broadly support Starmer's vision for AI in public services. Amanda Brock, the CEO of OpenUK, an open-source technology advocate, says: "Our governments have no choice but to bring in AI, but they need to do this in an informed manner."

The best response, she says, is to adapt to these changes, not fight them. "AI will be important to the public sector and we simply can't let it lag behind enterprise. The government is right to explore its use."

Bill Conner, chief executive at Jitterbit, an automation company, and a former GCHQ adviser, says enterprise AI can introduce "untold possibilities" for the government. But he warns a "rip-and-replace" approach to implementation could hinder internal adoption and introduce unnecessary risks.

"AI isn't a one-size-fits-all solution," he says. "Its success hinges on the humans who operate and manage it. Even with AI automating tasks, human expertise remains crucial for decision-making."

The government's commitment to workplace automation has reignited discussions in the public and private sectors about replacing human workers with machines. What do organisations hope to achieve by doing so? How should the process be managed? And how can its effectiveness be measured?

"If the goal is pure replacement, then the primary outcome will be cost reduction," says Dr Alexandra Dobra-Kiel, innovation and strategy director at Behave, a consultancy. "But cost reduction is questionable when considering the broader economic and social costs. A workforce displaced by AI doesn't just vanish, it creates ripple effects from mental health crises to shifts in tax revenues and public services."

Organisations will struggle to fully replace human employees with AI systems. Take coding, for instance – a task that GenAI excels at. Although AI systems can be useful in writing code, software engineers must still collate, document and audit what it writes. Even the most advanced agentic AI will require some human oversight.

Fabian Braesemann is a social-data scientist and the departmental research lecturer in AI and work at the Oxford Internet Institute. He co-authored a paper, *Winners and Losers of Generative AI in the Freelance Job Market*, which outlined the complex ways that GenAI tools such as ChatGPT are reforming labour markets.

While Braesemann is positive about the use of AI in the public sector, he notes that job roles consist of many different tasks, each of which requires different skills.

Naturally, some roles have been automated out of existence and few would suggest bringing them back. But for many other roles, outright replacement is not yet possible.

Braesemann gives the example of a painter and decorator. Most of their work could not be easily automated, but some cognitive tasks, such as writing and chasing invoices, could be. And many jobs like this, he adds. Attempts at automation therefore should prioritise replacing tasks, not roles.

Leaders must also understand that sometimes the use of AI systems can complicate workflows or introduce new challenges. For instance, because GenAI tools are susceptible to hallucinations, employees must spend time reviewing their outputs and correcting any errors.

Dobra-Kiel advises firms to evaluate AI against benchmarks for accuracy and error rates – the number of errors found in AI-generated decisions versus human decisions – to ensure these systems are not introducing biases or systemic flaws. Or benchmark cost savings against service quality, she says. "If AI reduces costs but degrades services, the trade-off may not be justified."

A complete measure of AI's impact on workplace efficiency and productivity would consider both speed and quality, says Dobra-Kiel. "AI's primary advantage is speed, so organisations should measure how much time is saved on a task without sacrificing quality."

Crucially, she adds, organisations must prioritise the ethical use of AI. This means, for instance, auditing the systems' outputs for demographic biases or measuring employee engagement before and after AI integration.

Ultimately, if the government struggles to integrate AI in the civil service effectively, it may be thanks to a lack of agility, not an absence of best practices. Unlike the digital-native businesses in the private sector, which are racing to adopt AI, public sector institutions are historically risk-averse and tend to implement organisational or operational changes slowly.

While Starmer may be keen to reap the benefits of AI in the public sector, this may create some organisational challenges when establishing more agile ways of working. ●

> A workforce displaced by AI doesn't just vanish, it creates ripple effects from mental health crises to shifts in tax revenues

BOTS WILL HAVE MORE WORK RESPONSIBILITIES THAN HUMANS BY 2030

Share of workplace tasks to be handled by machines and humans

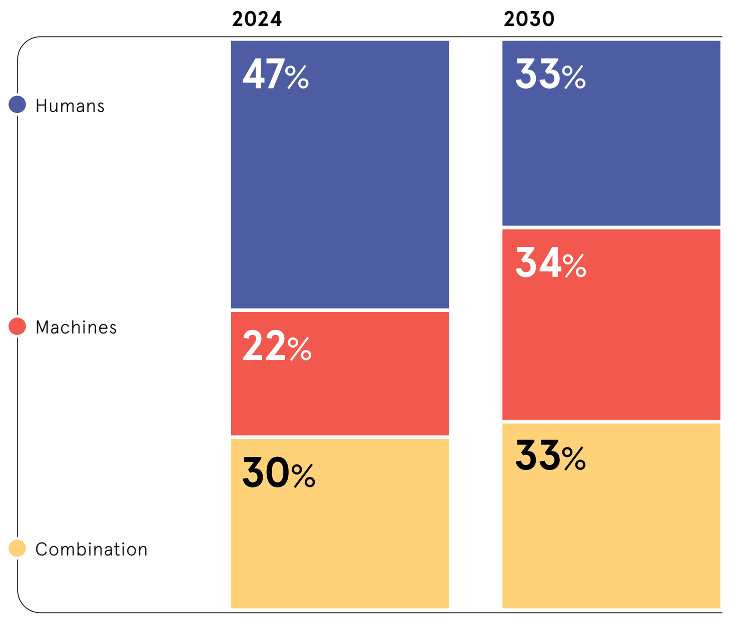

World Economic Forum, 2025

INTERVIEW

'We simply need better tech'

Carissa Véliz, an author, philosopher and ethicist, discusses the threat of a post-privacy internet and calls on businesses to stand up to the tech companies that are using their data to train AI systems

Tamlin Magee

Carissa Véliz, an associate professor of philosophy and ethics at the University of Oxford, says the age of AI presents new challenges in digital privacy. Véliz is a staunch advocate for privacy rights and the author of Privacy is Power, which outlines how governments and big tech companies track our digital footprints.

Here, she explains how privacy is being eroded in the generative AI era – and what companies should do to address it.

Q Opting out of some AI services is difficult but possible. But what about the systems that train AI models – can normal users opt out of those too?

A No, it's impossible to opt out, and that is a huge problem because it means that these systems do not respect privacy laws, whether in the UK or Europe, and yet they're still allowed to function.

I'm worried that, instead of forcing tech companies to follow the law, we will change the law to adapt to the technology, because the inconsistency threatens our rule of law. That is the wrong way to resolve this tension. Laws must lead, not follow, the technology.

There are two ways in which we are clearly not opting out. First, these systems are being trained on our data and they are using all data available on the internet: social media, forums and so on. And although we have the right to ask companies to delete our data, they don't even know what data they use.

Second, the only way to undo the data collection that's occurred so far is to delete the models that have been trained on our data and retrain them on different data – and that is not going to happen.

Q Should societies as a whole consider putting the brakes on GenAI?

A I don't like that kind of talk, because it creates defensiveness – we simply need better tech. Tech can be designed better to support democracy and privacy rights. We shouldn't put the brakes on it, we should make it better.

When you're using an AI chatbot such as ChatGPT, it's making a lot of inferences about you from the way you use language; for example, where you might live or what age you might be.

It's difficult to understand how much data you might be losing. It's not only what you type on your keyboard, but also what can be inferred from that – and you might not realise how much can be inferred.

Q Are current data and privacy regulations ill-suited for the GenAI age?

A I don't think it's fair to blame regulation. Part of the reason the rules are very far from perfect is that tech companies lobbied hard for soft regulations. The GDPR, for instance, would have been much stronger if companies hadn't pushed back. So we have this crazy system of rejecting or accepting cookies all the time thanks in part to efforts by the companies using our data.

It would be much simpler if firms didn't collect your data by default. Then you wouldn't have to say no to cookies every time you go to the website and you'd just have to say yes once.

Q There's a similar discussion happening now in the UK, which appears to be moving to an opt-out model for AI and copyright.

A It's crazy. It's the result of pressure from companies and it's not going to work because it puts the burden on individuals and that's not fair. But regulation is not at fault; the problem is that regulation is not going far enough.

Two years ago, regulation seemed to be strengthening. For instance, Europe was considering an enhanced privacy directive to fix some of the faults of the GDPR. Now we're going backwards. Thanks to US pressure and geopolitical tensions, regulating these things is becoming harder and harder.

> "The only way to undo the data collection that's occurred so far is to delete the models that have been trained on our data – and that is not going to happen

Q Businesses are also unhappy that their data is being used to train GenAI platforms – could we see pushback from the that community as well?

A Not only a pushback – *The New York Times*, of course, is suing OpenAI – but my hope is some companies will rise to the challenge. Take Proton, for example, which offers an encrypted email and productivity software suite. Full disclosure: I'm on the board of the Proton Foundation, but that's because I actually believe in them and I've been tracking them since they started. There are other similar services out there too.

Once companies up their standards, everyone else follows, including legislators and regulators. Companies can be the good citizens that improve standards for everyone. But when it comes to pushback, I really hope the newspapers stand their ground, because they're most at risk. If they capitulate to tech companies, I'm not sure how it's going to end.

Careless People, a new book about Facebook and its founder, contains a quote from Mark Zuckerberg predicting the end of newspapers. He essentially says there are two options: he can buy them or he can create his own. And he doesn't seem to realise, or care, how catastrophic that would be for democracy.

Q Digital sovereignty is becoming a serious political issue in Europe. Could the same happen with digital privacy?

A I expect so. We are already seeing it in the US, with worries about TikTok. There are two concerns here: privacy, meaning issues related to the apps, and control over algorithms, especially as the latter is what enables you to sway public opinion.

When you consider risks more broadly it becomes obvious why privacy is important. We should have privacy all the time, because you never know what's going to be a risk. Often, by the time you recognise the danger, it's too late. That's precisely the point of privacy: to prevent abuses of power.

Q You've said your students are increasingly avoiding digital platforms. Will we reach a point where people think twice before using AI systems or similar tech?

A We might reach an inflection point, yes. This is a huge battle because, although people are tired of being exploited by big tech, they're also incredibly busy and often feel overwhelmed in their professional and personal lives. And many digital platforms do make our lives easier, so naturally people succumb to convenience.

It will be a constant struggle to balance those competing factors: privacy and convenience. That's why we need to find convenient ways to preserve privacy. ●

Commercial feature

From compliance to control: mastering AI and data sovereignty

With top-tier infrastructure and strong policies ready, UK businesses must fast-track sovereign AI and data strategies to secure control, fuel innovation and stay competitive on the world stage

The global economy is entering an unprecedented phase of transformation, driven by the rapid rise of data and artificial intelligence.

According to a report by Forrester, by 2028 the global digital economy will reach a staggering $16.5tn (£12.2tn), making it the third-largest economy on the planet, behind only the US and China.

Meanwhile, the International Monetary Fund (IMF) forecasts that AI alone will drive 7% of global GDP growth over the next five years, more than double the expected growth rate of 3.4% for the broader economy. This shift is a fundamental reordering of economic priorities and competitive advantage.

The critical question for every organisation is clear: where will your growth come from in this new data-driven world?

"Data and AI are no longer optional tools or experimental technology, they have become the cornerstone of economic growth and the decisive edge in global competition", says Kevin Dallas, CEO of EDB, a leading enterprise data and AI platform provider.

Despite this urgency, EDB's global research, involving over 2,000 executives across North America, EMEA and APJ, reveals that only 23% of enterprises are actively building their own sovereign AI and data platforms.

These pioneers are pulling ahead – investing in sovereignty, observability and AI readiness to build platforms for autonomous, real-time decision-making.

At the heart of this movement is sovereignty: the ability to exercise full control over AI and data assets without sacrificing agility or compliance. It's a comprehensive approach that covers access, visibility and the ability to use AI and data when needed most.

"Data and AI sovereignty isn't about hiding behind a firewall or retreating from global collaboration," Dallas explains.

"It's about freedom – the freedom to choose your AI models, to keep data compliant with evolving regulations and to deploy capabilities across clouds, borders and teams without compromise."

According to EDB's research, 97% of enterprise leaders see becoming their own AI and data platform as mission-critical, yet only 63% understand that sovereignty is essential to achieve it.

Without it, organisations risk agility without control, leading to fragmentation and lost opportunities.

"Building an AI and data platform isn't simply about technology procurement, it means bringing every tool, model and dataset into one secure, extensible environment where they can operate seamlessly together," Dallas points out.

The foundational technology enabling this is evolving. Solutions like Postgres offer a single architecture capable of handling structured and unstructured data, supporting transactional, analytical and AI workloads alike. This versatility is essential as enterprises move from experimentation to scaling production AI.

Those already leading the way have begun building what Dallas calls "agentic AI factories." These are internal AI ecosystems designed to deliver hyper-personalised services and autonomous outcomes across multiple business domains.

According to EDB's research, the 13% of organisations investing heavily in such systems report nearly three times the expected ROI compared to peers.

In highly regulated industries – financial services, healthcare, defence and public sector – the pressure to scale agentic AI securely is intense.

"A sovereign platform that is hybrid by design makes this possible. It gives organisations the flexibility to run AI where their data resides – be it on-premises, across multiple clouds, or at the edge – while maintaining full observability and control over the entire data estate," says Dallas

This approach safeguards sensitive information and ensures regulatory compliance without stifling innovation.

Currently, just under one in four enterprises globally understand this urgency. But projections show that, within three years, half of all organisations will recognise sovereignty and AI readiness as mission-critical. This is a short window – one that demands swift strategic action.

Success will require hybrid deployments that tightly couple data and AI, ensuring both are secure in motion and at rest.

AI systems must be flexible, safe and production-ready. And, importantly, the underlying platforms must be open and extensible – not confined by proprietary technologies or legacy constraints.

"This is about more than competitive advantage," Dallas stresses. "It's about national and economic resilience. The UK has the talent, infrastructure and policy momentum. What it needs now is the commercial will to turn that potential into real platforms and capabilities."

For UK businesses, the risks of delay are clear. Falling behind in sovereignty and AI readiness threatens exclusion from emerging value chains, regulatory fines and a loss of customer trust.

As sovereignty becomes a key differentiator, companies relying heavily on third-party platforms could risk reputational damage and diminished investor confidence.

The UK's National AI Strategy has laid important groundwork – committing to secure, explainable and trustworthy AI ecosystems. Investments in computing infrastructure, including the AI Research Resource and Isambard-AI supercomputer, are among Europe's most significant. But government efforts can only pave the way; enterprises must take the wheel.

> "With the right platform architecture, organisations don't have to choose between openness and control – they can have both

78%

of business leaders say genAI has increased their attack surface

72%

say they have increased their risk management in AI governance

PwC, 2024

"Government can build the roads, but businesses have to drive the cars," Dallas remarks. "That means embedding sovereign AI and data governance into your core digital strategy, investing in talent and committing to platform ownership from day one."

Deploying AI responsibly is not simply about capability but accountability. Sovereign AI ensures compliance, aligns with business goals and allows organisations to innovate with confidence and transparency.

Ultimately, sovereignty is not about isolation. It can enable global interoperability, adaptability and resilience, equipping organisations to compete confidently in a complex, evolving regulatory landscape. From GDPR in Europe to data localisation in Asia and cloud compliance in the US, the ability to adjust systems dynamically is critical.

"Flexibility built on control is the new foundation," Dallas concludes. "With the right platform architecture, organisations don't have to choose between openness and control – they can have both."

The competition for influence in the global AI economy is intensifying. Sovereign readiness will determine who captures the most value as digital transformation accelerates.

"There is a narrow window for the UK to assert itself," Dallas warns. "Every day counts. Those who transform intent into execution today will lead the next thirty years of growth."

The question now is whether UK enterprises are ready to make data and AI sovereignty their strategy before the window closes.

For more information please visit enterprisedb.com

EDB POSTGRES AI

oneAdvanced

Your work powered by Ai

Safe. Trusted. Secure.

Our software features sector-specific AI agents to streamline tasks, UK-based data processing for security, and customisable privacy controls to ensure compliance. Allowing you to innovate without compromise.

Find out more at: www.oneadvanced.com/ai