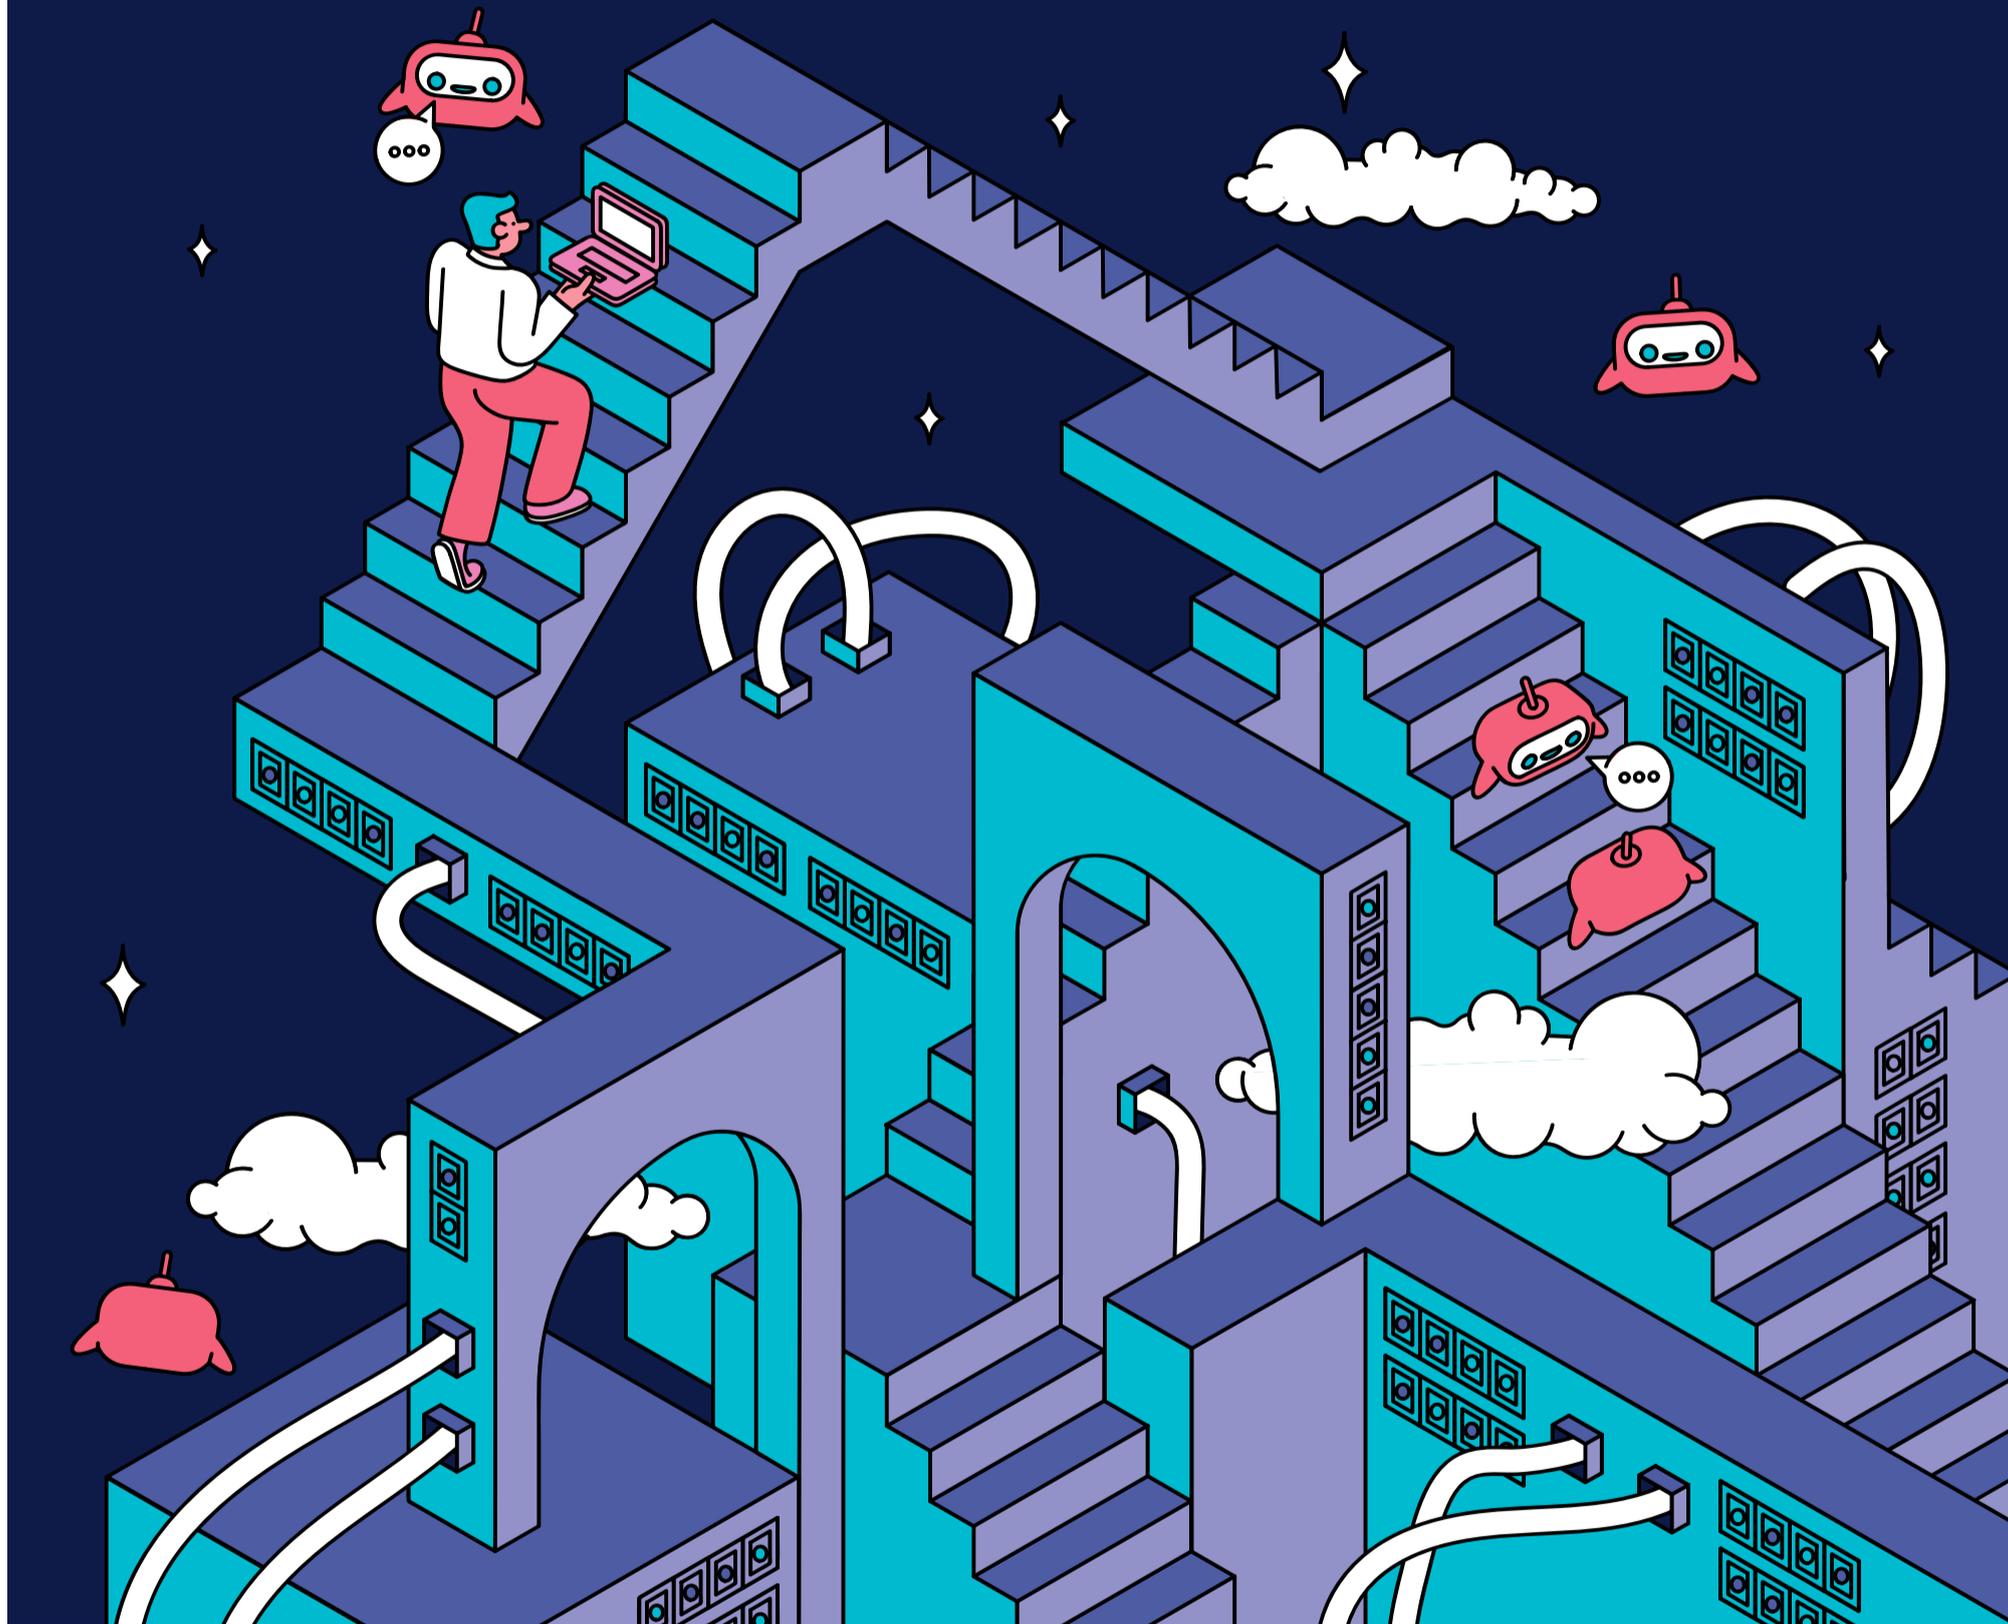


ENTERPRISE AI

09 FIVE EXPERTS WEIGH IN ON THE EU AI ACT

12 MAPPING THE MOMENTS THAT MADE AI IN 2025

16 CAN WRITERS TRAIN AI TO MIMIC THEIR STYLE?



telana

experts in
applied
innovation



SUPERMICRO Accelerate Enterprise AI

First-To-Market Solutions Simplifying At-Scale AI Deployments

Learn More at
www.supermicro.com/EnterpriseAI



ENTERPRISE AI

Distributed in
THE TIMES

Contributors

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

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
Jessica Lynn

Commercial production managers
Alex Datu
Audrey Davy
Ellen Newsome

Design and illustration
Kellie Jerrard
James Lampard
Celina Lucey
Samuele Motta

Design director
Tim Whitlock

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. The information contained in this publication has been obtained from sources the Proprietors believe to be correct. However, no legal liability can be accepted for any errors. No part of this publication may be reproduced without the prior consent of the Publisher. © Raconteur Media

in [raconteur-media](https://www.raconteur-media.com) @[@raconteur.stories](https://twitter.com/raconteur.stories)

raconteur.net /enterprise-ai-september-2025

SKILLS

UK faces upskill battle

The government has unveiled an ambitious nationwide programme to upskill workers in AI. Implementing it will be complicated work

Tamlin Magee

Keir Starmer has positioned AI at the centre of the government's industrial strategy, with the aim of making the UK an "AI superpower". To achieve that goal, the government has organised a nationwide skills drive to develop home-grown AI talent.

It's an ambitious plan. Schools, colleges, community hubs and workplaces across the UK will soon offer AI-skills courses to all citizens. The aim is to equip 7.5 million people – about one-fifth of the workforce – with the skills to use generative AI effectively.

Hyperscalers including Microsoft, Amazon, Google and IBM have teamed with the government to advise on the creation of AI training modules. Technologists, too, have praised the initiative.

Alexandra Dobra-Kiel, innovation and strategy director at Behave, a consultancy, says the government's efforts are timely and necessary. "Equipping people with the skills to engage with, rather than be displaced by, AI is crucial and the collaboration with major tech firms shows a serious commitment to preparing the workforce for a rapidly changing economy."

Nearly 3 million people in the UK have no access to the internet or digital devices, according to figures from Ofcom. What's more, research suggests that half of working-age adults cannot complete all 20 of the foundational digital tasks deemed essential for modern work, and nearly 8 million UK adults have no basic tech skills at all. So is the government perhaps putting the cart before the horse?

Engagement with AI tools varies significantly across the population. To successfully upskill such a large share of citizens, providers must tailor their training modules appropriately, ensuring everything from foundational principles of AI to advanced uses is covered.

That's according to Richard Giblin, head of public sector and defence at SolarWinds, an IT company. "Core digital literacy, coding and data skills must be in place before we can meaningfully scale up AI expertise across the workforce." Once those fundamentals are mastered, he says, organisations should take a layered approach to AI upskilling. At the most basic level, every professional must understand what AI can and can't do, as well as its practical applications and ethical implications.

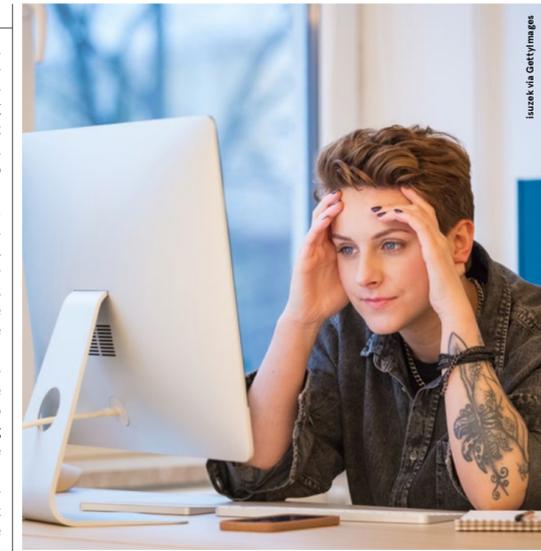


Image via Getty Images

Although GenAI is constantly improving, it remains prone to wild hallucinations and biases. And, because these systems often work to merely validate users' hypotheses, in the most extreme cases, GenAI platforms have encouraged delusions among their users.

Those seeking to use the tech effectively must first understand how GenAI works, says Stephan Reiff-Marganec, head of computing and engineering at the University of Derby. AI providers are all too happy to let magical thinking swirl around their software. Some have even issued premature statements about machine sentience. But, he explains, GenAI tools are really "just clever pattern-matching [machines]".

Nationwide upskilling initiatives should focus on helping users understand "how AI tools work and what's sitting behind them – how they're using the data that you're putting in and how you can control that", Reiff-Marganec says. This

means training people to use AI ethically and interrogate its outputs. Doing so could also help to boost critical-thinking skills across the population, an additional benefit of any AI-upskilling initiative.

Only when these basics have been covered should courses proceed to advanced topics, such as prompt engineering, coding or industry-specific skills. Here, too, learners should be encouraged to limit their use of the technology to tasks it is best suited to handle, such as data collection or search.

According to Reiff-Marganec, AI isn't a shortcut, it's a "tool that can guide you and give you ideas that you might want to look at. But you need to add your own understanding to push it in the right direction."

Giblin agrees, adding that any organisation seeking measurable benefits from AI adoption must view the technology holistically and avoid treating it as a shiny new tool. "AI can't operate in isolation,"

he says. "We will still need IT professionals to provide the infrastructure, cybersecurity and oversight that underpin AI adoption, checking for accuracy, compliance and long-term sustainability. Building AI skills must go hand-in-hand with strengthening complementary IT disciplines."

The government is right to prioritise AI, he says, but foundational skills must not be overlooked. AI training should focus on ethical and practical uses of the technology. Only then can we really "future-proof the workforce", he argues.

Any government-led upskilling initiative should prioritise the public good over commercial interests. Dobra-Kiel acknowledges that big-tech firms should be included in the government's efforts, but she cautions that their role must be to contribute to the curriculum, not to control it.

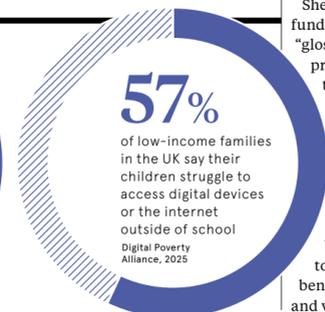
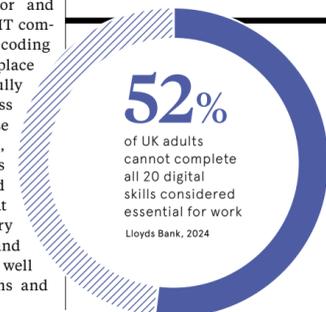
"Big tech can provide valuable insights on real-world tools and rapidly evolving technologies," she says. "Their involvement is essential. But allowing them to dictate the curriculum risks embedding commercial interests, tool-specific training and vendor lock-in, teaching people about products rather than principles. The curriculum should reflect public interest, not just corporate roadmaps."

Because technology evolves so quickly, traditional training modules will inevitably fail to keep pace with the subject matter, Dobra-Kiel explains. Courses should, therefore, aim to help people develop durable learning habits rather than "ticking off toolkits".

Any initiatives must cut across functions and disciplines. "Treating AI as a purely technical subject is a mistake because AI touches law, ethics, design, politics and power," says Dobra-Kiel. "We don't just need more prompt engineers, we need teachers, careworkers and shop-floor staff who know how to question an algorithm's bias or understand data privacy rights."

She emphasises that publicly funded training must not stop at "glossy success stories". Training providers should ask whether they are building capabilities or merely credentialing. "Accountability requires opening the black box of training outcomes, not just celebrating intentions."

That might mean examining some uncomfortable data, but doing so can help educators to better understand who is benefiting from the training – and who's being left behind. ●



Agentic systems pave the way to AI autonomy

The shift from AI pilots to autonomy is underway, as businesses begin to embed agentic systems into workflows

Most organisations are still adapting to AI. Initial experiments with the technology have focused on chatbots, pilots and small-scale proofs of concept. While these were useful for building awareness, they rarely shifted the dial on business performance in a significant way – until now.

For most firms, the question is no longer whether they should adopt AI, but how deeply it should be embedded in core operations. This is why agentic AI has become the business buzzword of the year. But unlike many trends, this one has the potential to be a genuine differentiator for organisations.

Agentic systems move beyond narrow automation to create autonomous workflows. They don't just execute tasks, they set steps, adapt dynamically and are capable of achieving outcomes with minimal human input. Gartner forecasts that by 2029, agentic AI will autonomously resolve 80% of common customer service issues. That level of independence marks the dawn of AI autonomy, where agents become collaborators rather than tools.

"We've moved from that chatbot phase of AI into thinking about how we embed it into our core workflows," explains Amanda Grant, chief product officer at OneAdvanced. "The differentiator now isn't whether you use AI, but how deeply and intelligently it's integrated into your operating model."

For business leaders in the middle of the AI journey, agentic AI represents the critical inflection point. Getting it right means shifting from proof-of-concepts to structured workflows, from generic tools to sector-aware agents, from cost savings to customer experience and from fear of risk to confidence in trust.



Organisations that just experiment with AI without a focused plan will probably fail to achieve their goals

Here, we lay out how agentic AI is moving from hype to practical business advantage, with six key shifts that business decision-makers should pay attention to.

The shift from experimentation to real value

The last wave of AI pilots helped test what was possible but often sat on the edge of the business. Now, organisations are embedding AI directly into workflows. Instead of a chatbot in a corner, businesses want systems that operate "in the flow of work".

That change requires integration with data, applications and decision-making processes. The focus is on AI's measurable value.

With expertise in healthcare, government and legal services, OneAdvanced helps users design AI-enabled processes that complete real-world, business-critical jobs rather than experimental side projects.

"Organisations that just experiment with AI without a focused plan will probably fail to achieve their goals," says Grant. "Real progress comes when you understand the critical jobs to be

done – summarising a medical document, progressing a legal case, processing a complaint and so on – and build AI into that flow of work."

Agentic AI as a business differentiator

An AI agent handles a single task, such as summarising a document. An agentic system goes further: it identifies the problem, plans a sequence of steps and executes them autonomously.

That means processes can run continuously without waiting for human intervention. A single engineer could orchestrate a team of AI coders – one writing code, another testing it, a third checking architecture and so on. Humans become the conductors, not the performers.

"Agentic AI is a huge productivity gain because you can scale up without bottlenecks. Agents can work overnight while your team rests. That makes them genuine business differentiators," says Grant.

Crucially, the advantage doesn't come from generic tools. Success depends on sector-specific AI agents that understand industry knowledge, regulations and workflows.

From cost savings to customer impact

Early AI adoption was framed in terms of efficiency and cost-cutting. While savings still matter, the new frontier is customer experience.

Customers expect faster resolution, personalised interactions and predictive outcomes. Agentic AI, with its ability to anticipate needs and take action, sets a new standard.

Gartner's forecast that eight out of 10 customer service issues will be resolved autonomously in the not-too-distant-future implies not just cost savings but a fundamental shift in expectations. The dream is that waiting in a queue for a human agent may soon be a thing of the past.

According to Grant, the conversation has moved from cost efficiency to customer value. For example, in healthcare, the goal isn't just reducing admin time but improving patient outcomes by giving clinicians better information and more time with patients.

The importance of trust and security

Autonomy brings both opportunity and risk. For instance, an agent instructed to maximise performance might spin up a prohibitively expensive server unless boundaries are set. That's why guardrails are non-negotiable.

"The more autonomy you give, the more you must have these rules in place. With agentic AI, trust and security step up a notch," says Grant.

Every agent should operate against well-defined goals. At the same time, there should be clear limits on what the AI can and cannot decide on its own. At critical junctures, humans provide the judgment and accountability that machines cannot replicate.

Security also extends to the data itself. Controls, such as detecting and masking personally identifiable information, protect sensitive inputs and outputs from misuse. Auditability is equally important. Every decision made by an agent must be traceable, with transparency of the reasoning process and the data that informed it.

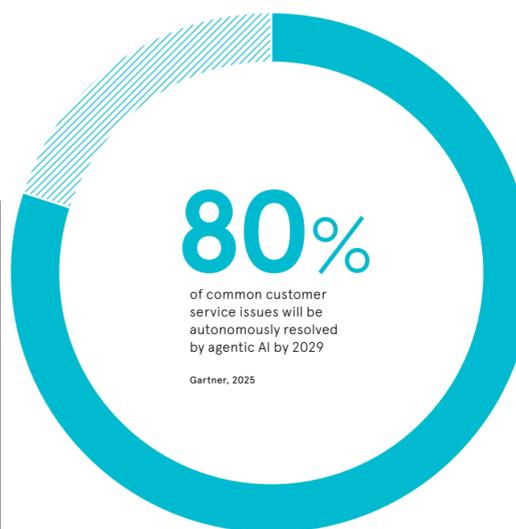
Compliance, too, is essential to ensure that systems remain lawful as well as trustworthy. Finally, strong cyber defences must underpin the entire stack, from the models themselves through orchestration layers to continuous monitoring.

By combining risk-based permissions with transparency and auditability, organisations can grow without losing control.

Real-world proof, not hype

Agentic AI is already delivering results. Take UK GP surgeries, where OneAdvanced has deployed its Workflow Assist systems.

Each surgery receives hundreds of electronic documents per week, ranging from specialist letters to hospital updates. Staff must summarise them, triage urgency, identify medication



changes and assign SNOMED clinical codes – universal reference numbers that translate written medical details such as symptoms, diagnoses and procedures into a standardised code that both humans and computers can understand – from a library of 360,000 terms. The process is time-consuming, can be prone to error and adds to NHS staff pressures.

OneAdvanced's system uses two AI agents: one summarises documents, flagging urgency and next steps, while another suggests relevant SNOMED codes.

But humans still validate outputs and 90% of user feedback has been positive on accuracy. With estimated average time value savings of £12,000 per practice per year, and even greater improvements in patient care, the solution is proving its worth.

"What matters is not just saving admin time," says Grant. "It's the better outcomes – if GPs get clearer summaries and accurate coding, patients get better care. That's the real value of agentic AI." See the detailed case study [right] for more on this deployment.

Preparing for the future of work

Headlines about "the end of the entirely human workforce" miss the point. The future is not human versus AI, but human with AI.

Agentic systems will handle repetitive and predictable work, freeing people to focus on creativity, strategy and complex decision-making. Roles will evolve, where an engineer can become the orchestrator of a team of digital coders, not the sole producer.

Grant explains: "Organisations must focus AI on the jobs to be done, embed it in the flow of work and build trust through transparency. That's how businesses will unlock the real promise of AI autonomy."

Beyond the pilot phase: scaling with confidence

What distinguishes organisations making real progress with agentic AI is not experimentation but scale.

Moving from pilots to production requires confidence that the system



Agentic systems will handle repetitive and predictable work, freeing people to focus on creativity, strategy and complex decision-making

can deliver consistently, safely and in compliance with regulation. That's why the focus is increasingly on auditability, trust and sector expertise. Leaders want to know not just that an agent can perform a task once, but that it can repeat it thousands of times across complex workflows, without drift or unintended consequences.

Equally important is the change management that accompanies deployment. Agentic AI isn't just a technology upgrade – it reshapes roles, responsibilities and even measures of performance. Successful organisations invest as much in training and culture as in algorithms. They recognise that humans remain accountable and that transparency is essential to maintain trust with employees, regulators and customers.

By moving from AI experiments to purposeful agentic solutions, firms can unlock the transformative potential of the technology.



For more information please visit oneadvanced.com

oneAdvanced



How agentic AI is transforming GP surgeries

The challenge

The NHS faces a perfect storm: constrained funding and staff shortages set against a backdrop of rising demand. GP surgeries are on the front-line of this pressure.

Every day, surgeries process up to 600,000 documents nationwide – referral letters, hospital discharge notes and more. Each document must be read, assessed, summarised, coded and filed into patient records.

With 360,000 SNOMED clinical codes to choose from, accuracy matters. Coding errors can impact patient care and even funding, as certain codes drive resource allocation. Yet the manual workload drains staff time and risks burnout.

The solution

OneAdvanced worked with GP surgeries to co-develop the Workflow Assist system. It includes two AI agents: one that processes and condenses multi-page letters into clear summaries, highlighting urgency, medication changes and next steps, and one that suggests the correct SNOMED clinical codes to align records and ensure consistency.

A human remains in the loop to validate outputs, ensuring safety and building trust. Over time, more autonomy will be added to the AI.

"At the moment we've got humans validating, but the next stage is agentic – where the system doesn't just suggest but also acts, like auto-filing non-urgent letters," says Grant. "But it's all about trust. That's why auditability and transparency are essential."

The impact

Workflow Assist has been transformative. Users report a 90% approval rating for the accuracy of processed documents. On average, each practice can save £12,000, largely thanks to reduced time spent on administrative tasks. Patients benefit, too, with GPs receiving clearer information more quickly, which leads to better outcomes. Notably, 35% of new customers now come through referrals, reflecting strong trust and satisfaction.

Grant says: "The cost savings are significant at a national scale, but the bigger impact is freeing up GPs to spend more time with patients – potential capacity to deliver 150,000 more patient appointments a week. That's what matters most."

Why it works

The deployment succeeds because it meets three key criteria. First, it focuses on clearly defined jobs to be done – summarisation, coding and filing – each with measurable impact.

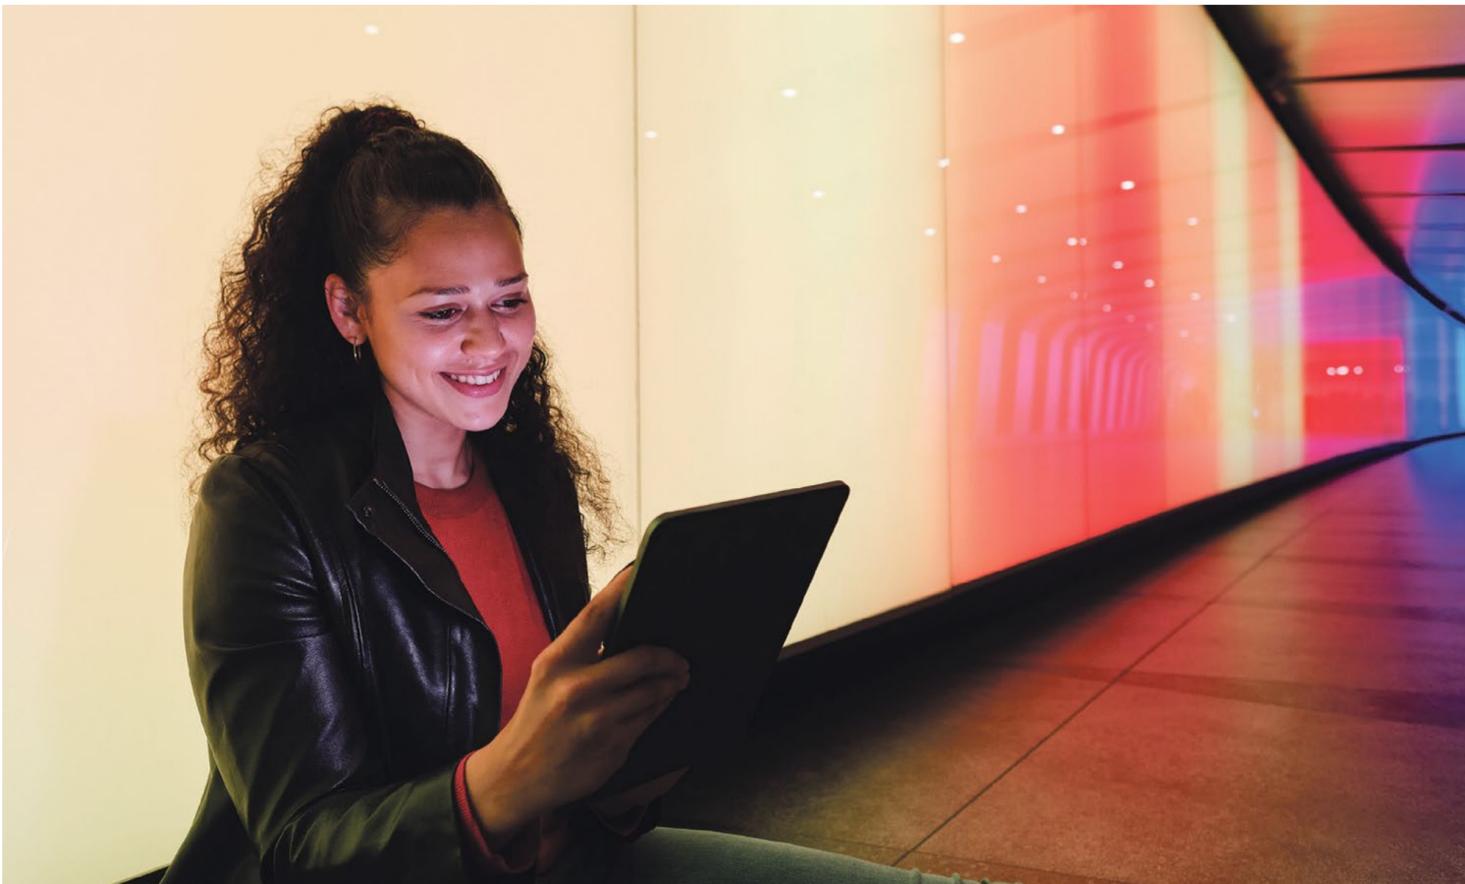
Second, clinicians remain in the loop, retaining oversight that ensures trust in the system.

Third, the solution reflects OneAdvanced's understanding of GP workflows, regulation and funding models. Unlike AI for the sake of it, this system delivers real-world results precisely because it is tailored to the context.

The future

The roadmap includes adding greater autonomy, such as automatic filing of non-urgent documents and workflow orchestration. Each step will balance increased automation with the safeguards of auditability and transparency.

Ultimately, this case demonstrates how agentic AI can move from hype to genuine transformation – not by replacing humans but by enabling them to focus on higher-value work.



FINANCE

5 ways AI is transforming the finance function

Once reluctant to entrust high-value work to AI, CFOs may finally be warming to the idea. Here are five ways the technology is reshaping financial operations – for better or worse

Sam Birchall

Finance leaders usually try to avoid using systems they don't fully understand – and rightly so. Mistakes in the books or improper financial reporting, for instance, can lead directly to big losses. But many, it seems, are finally coming round to implementing AI.

Nine in 10 finance chiefs now believe generative AI is crucial for financial planning and analysis, according to a new report by Pymnts Intelligence, a data platform. Fewer than half expressed such confidence in March 2024. The findings signal an increased reliance on GenAI across core financial

functions, including reporting, cost management and strategic planning.

But AI systems can do much more than crunch numbers and automate spreadsheets. From deciphering behavioural cues to decoding trader slang, AI's role in finance is expanding in new and unexpected ways.

Translating finance-speak

Financial criminals and crooked traders plan their activities in coded or jargon-laden exchanges when they know their telephone conversations and emails are being recorded. That's why compliance providers are developing large language models capable of decoding clandestine communications.

Banks, hedge funds and regulators can use these tools to track employee messages and uncover secret languages. One company claims its pilot system can decipher phrases coded with emojis and even pig Latin.

Meanwhile, boards are using GenAI to produce natural-language narratives that translate complex financial-performance data into clear, accessible summaries. This frees up finance leaders' and analysts' time and helps to ensure that reports are intelligible to those lacking technical finance expertise.



Boosting market resilience

Navigating volatile markets is a perennial challenge for finance leaders. Thankfully, new tools are emerging to help analysts and consultants devise and direct major business projects. By analysing vast data sets to find patterns in a company's operations, supply chain or customer base, AI systems can unlock value and reveal opportunities for diversification, before market pressures force a reaction.

Researchers at Nottingham Business School's Centre for Business

and Industry Transformation have developed a system that analyses the potential value of companies' competencies, rather than their current products, and suggests how their resources can be redeployed in new markets.

For instance, the system was used by an autoparts manufacturer to identify in-house proficiencies in metal processing. These insights enabled the firm to pivot into a new area of business: jewellery exports. Then, when the market reached saturation, the firm pivoted again to manufacturing metal wound clamps for battlefield medics.



Detecting emotion

Bankers and other finance professionals are using AI tools to enable better interactions with customers. AI providers have designed systems capable of understanding human emotions. These can detect when specific words, phrases or pauses might indicate feelings such as frustration or confusion and guide service agents to respond effectively.

In finance, customer service tactics rely heavily on demographic data, such as age and income. Emotional analytics, however, add a layer of immediacy, enabling banks or financial advisers to adjust their advice based on how a customer is feeling at that moment.

The application of these systems in finance extends far beyond customer care. In procurement or negotiation calls, for instance, analysing tone or speech patterns could help teams spot hesitation, tension or evasiveness. Similar tools are being developed to serve vulnerable customers. These flag vocal cues that might indicate dementia or a mental illness, helping finance professionals implement relevant safeguards.



The global emotion-detection and recognition market is forecast to reach \$136bn (£101bn) by 2030, according to Grand View Research, fuelled by innovations in deep-learning algorithms.

But the integration of emotional AI in financial services is not without its challenges. Chief among them are the ethical and privacy concerns that arise when firms begin deploying surveillance tech and manipulating customers' emotions. In the UK, a parliamentary inquiry on AI in financial services recently received written evidence that both lauded emotional and behavioural surveillance and warned of its harmful effects.



Predicting burnout

Burnout looms large in the finance function. Almost all (99%) accountants have experienced burnout resulting from relentless work pressures, according to a study by Floqast, an accounting-software firm.

Some companies are using AI to scan workplace data and flag signs of stress or overwork before the problems spiral. By analysing expense-claim patterns, for example, AI can spot the tell-tale markers of exhaustion, such as frequent late-night meals or a surge in last-minute travel bookings.

Unilever, for instance, has rolled out an AI-powered wellbeing programme that monitors workplace data for potential burnout risks. It has been integrated with the Microsoft 365 suite to analyse calendar and email patterns while maintaining strict privacy controls.

However, although companies argue that such technologies support employee wellbeing, in practice, these systems often amount to intrusive workplace surveillance. If not implemented carefully, AI-powered workforce monitoring software could harm rather than help workers' mental health and morale.

Hunting down fraud

Sam Altman warned that we are facing an impending AI-powered fraud crisis, where anyone will be able to perfectly imitate anyone else. This prediction has finance leaders worried. Last year, the British design firm Arup was scammed out of \$25m (£19m) by a deepfake. And such sophisticated attacks are becoming more common.

To combat this, some firms are trialling AI systems that go beyond flagging suspicious transactions and can analyse the behavioural biometrics of users.

During high-risk transactions, for example, these systems can detect if a user is typing erratically, hesitating, switching devices or appears to be coached over the phone. Such patterns can indicate social engineering or coercion in real time.



Some tools even include a daily adaptive model, which retrains itself every 24 hours based on new behavioural data and fraud signals. Systems equipped with this function can detect emerging threats without human reprogramming. ●

INSIGHT

'The UK must build a culture that rewards risk-taking'

The government's AI upskilling initiatives won't be enough to even the playing field for startups and small businesses

We should all support the UK's ambition to be a leader in AI. But the manner in which the government's AI upskilling strategy is being implemented threatens to put the cart before the horse and overlooks how innovation really works.

The government has unveiled nationwide upskilling initiatives aiming to deliver essential AI training to 7.5 million workers by 2030. But much of the focus in practice has gone to hyper-regional efforts, such as strengthening the Oxbridge growth corridor. This risks overlooking clusters of talent in other parts of the country and reinforcing the North-South divide.

AI talent doesn't just sit in labs or lecture halls. It's found in small businesses across every region, where people are experimenting, building and creating the products that will define the next decade. If upskilling efforts don't reach every region and industry, opportunities will be concentrated too narrowly.

The UK has always been strong on research. The challenge, historically, has been commercialisation. There are simply too many ideas that remain in the lab or get swallowed up by big companies, because startups aren't given proper support to grow. Policymakers risk prioritising large corporations over the UK's own AI ecosystem, which largely consists of small, incredibly talented teams trying to expand and develop their businesses. They require different types of support than international tech giants.

For most startups, the biggest obstacle is not a lack of ideas or ambition, it's the process of hiring talent. The Office for Budget Responsibility estimates the rises to employer National Insurance contributions will add a 7.1% real-term rise to the price of hiring a full-time employee on the national living wage and 10.6% for part-time workers. Even smaller gestures, such as introducing National Insurance contribution cuts tied to revenue or head count, might give startups more resources to hire and retain UK employees.

Immigration policy in the UK may also hold back progress. The number of science, research and engineering visas decreased by a third in the later months of 2024. This is a problem for the UK's AI readiness. Hiring the world's best AI talent doesn't mean opportunities for domestic workers will be lacking.

Solving a lot of these problems comes down to mindset: the UK is still too risk-averse. More often than not, failure is seen as career-ending rather than as a stepping stone to better things. The UK must build a culture that rewards entrepreneurship and risk-taking.

Grant schemes such as Innovate UK should be sufficient, but limited categories, time-sensitive deadlines and complex documents and hurdles mean smaller businesses often don't bother to apply. In many cases, they lack the resources or knowledge to do so. Large enterprises with dedicated grant teams scoop up those benefits instead.

The UK must invest in AI capabilities. However, without the right foundations, training programmes will fail to deliver lasting impact. Most companies are ready to adopt AI in their workflows, but they need better support to embed AI into their systems and operations. That means not just training people to write better prompts.

For the UK to be an entrepreneurial country where people are encouraged to innovate, it'll need focused investment, along with government policy that reflects the way innovation actually occurs.

This can only be achieved with more inclusive, national talent initiatives, wiser immigration and tax incentives to help startups hire and retain talent, funding systems that work for small businesses and a shift from focusing on prestige projects to enabling real-world impact across sectors and regions.

If we get this right, the UK can fulfill its ambition of being a world leader in AI. But if we get it wrong, we run the risk of training a generation of people who will take their ideas and ambitions elsewhere.



Muj Choudhury
Chief executive
Rocketphone.ai

Your board will ask about AI for finance.
This is the answer.



Generic AI is for emails. Specialised AI is for your financials. Introducing Verity by BlackLine.

Built on two decades of financial automation leadership, Verity by BlackLine provides the auditable accuracy your business demands. It's the confidence you need for unstoppable finance.

 **BLACKLINE**





AI-native marks the end of 'app for that' era

Most businesses use AI, but only those that embed it at their core through an AI-native approach can truly learn, adapt and thrive

Business leaders have long been tasked with maximising the potential of human intelligence. But in 2025, they're not just managing the hearts and minds of people. Leaders must exploit the competitive advantages of AI and empower their human employees to use the technology effectively to steal a march on their rivals.

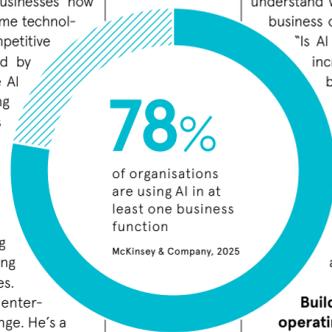
The majority of businesses are using AI in some capacity. A survey by McKinsey & Company found that 78% of organisations are using the technology in at least one business function. But early experimentation has led to fragmented, disconnected tools in isolation across different areas of the business. And while this might have been a necessary entry point into enterprise AI, companies must evolve their approach as the market and technology mature.

Isolated tools, such as AI-powered customer-service chatbots, are designed to solve specific, niche problems, rather than drive enterprise-wide transformation. This means they have little or no reuse potential in different organisational contexts, limiting their ROI. They also create complexity in the form of siloed data, technical debt caused by using multiple tools and models, and potential security risks in the absence of transparency and controls.

AI tools v an AI-native approach
Simply using AI tools won't be enough to maintain a competitive edge in the long term. After all, most businesses now have access to all the same technology. Instead, the competitive advantage will be seized by businesses that become AI native. This means moving away from treating AI as just another product or a bolt-on tool for isolated problems. Instead, businesses should embed it at the core of their operations, decision-making and processes for building new products and services.

John Clarke is helping enterprises navigate this change. He's a senior technology consultant at Telana, a leading provider of AI, data and cloud-technology solutions.

“The richer the data sources from across the organisation, the more insights and different questions you can ask of it



The company is helping businesses harness the full potential of AI technology to solve business challenges. Clarke says leaders must begin by adopting an AI-native mindset. This involves a whole organisational view, including building platforms, not just products, using data effectively, upskilling the workforce and adopting a continuous learning approach.

"Leaders are enthusiastic about using AI, but a technology-led approach has resulted in a lot of failed uses," says Clarke. "If you're designing a new product, service or process, leaders must understand whether there is actually a business case to solve," says Clarke.

"Is AI the right solution? Will it increase revenue? Will it drive business-wide change? How are you going to measure that? What are the metrics? My job is to get into that detail with leaders and measure the potential impact of an AI solution on operational efficiency, risk and revenue."

Building an AI operating system

Once those questions have been answered, the next step is to create the necessary infrastructure to facilitate an AI-native transformation. This involves building an AI operating system (OS) that serves as a central hub and nervous system for a business. The aim of an AI OS is to ingest, unify and extract insights from data across a company by considering workflows that connect all departments and integrating key platforms, including for customer relationship management and enterprise resource planning.

This enables businesses to weave AI into the fabric of the entire business: the

products and services delivered to customers, the internal processes that drive day-to-day efficiency, the skills and mindsets of the people who use it, and the culture that prioritises data-driven, AI-informed decisions. In short, the OS would create an intelligent engine for running a business where AI is the default way an organisation develops.

Clarke says a central OS would generate intelligent and profitable insights that disconnected, siloed systems simply can't provide. "The ultimate aim is to build platforms, agents and a repository of data built through the AI models," he says. "It can start to answer multiple types of questions that you might want to ask about that data. The richer the data sources from across the organisation, the more insights and different questions you can ask of it. As a business acquires more data, it will continue to learn, adapt and evolve."

AI-native businesses in action

Telana is helping traditional businesses such as Channel 4 operate like AI natives. Partnering with the broadcaster, Telana is supporting Channel 4's ambition to help smaller brands create TV ads using generative AI. This initiative helps address a key market challenge: the high cost and complexity of producing ads that have historically been major barriers for smaller brands.

According to Telana, this AI-native approach enables these firms to punch above their weight. By using AI models to generate broadcast-quality ads, they are making TV advertising more accessible. The next step is to integrate AI into every stage of the creative process, from ideation and storyboarding to compliance checking and publication.

For AI-native startups, this is already their reality. This new generation of companies has been born with AI at its core. AI is the product, the primary workforce and often the key decision-maker. Algorithms are the brains that find market gaps and generate product ideas. They're also capable of writing and deploying code. Customer support and acquisition are also handled by AI, while financial decisions and transactions are carried out by AI agents. In these businesses, humans oversee tasks but don't perform them.

These startups provide a competitive warning to legacy organisations. Such companies learn at a much faster rate, with every transaction or click serving as a data point for continuous improvement. They're capable of making faster and more accurate decisions across key functions such as finance, operations and marketing. Predictive models also minimise risk by spotting forecasting errors and potential customer churn ahead of time.

Businesses that can adapt rapidly to this intelligent new reality will thrive in an AI-native future. Those that don't will be left behind by competitors who learn, adapt and grow at machine speed.

For more information please visit telana.com



LEGISLATION

Experts weigh in on EU AI Act

Brussels' landmark AI legislation is imperfect but could pave the way for similar policies around the world. What does it get right – and wrong?

Tamlin Magee

The EU AI Act is pioneering legislation that seeks to regulate transparency and accountability for AI systems and sets out acceptable levels of risk for AI uses based on societal, ethical and legal considerations.

EU member states will enforce the act, with penalties for the most egregious infringements reaching €35m or up to 7% of annual worldwide

turnover, whichever is higher. Similar to GDPR, the regulations apply to any organisation operating in the EU, not just businesses that are located in the bloc.

In August, a key provision, targeting general-purpose AI platforms, came into force. AI providers will be required to thoroughly assess the safety of their models through risk assessments and testing. The must

also maintain technical documentation about their model architectures and make those records available to authorities. They will also be required to publicly disclose their training data if it is requested by the appropriate authorities.

Here, experts in the fields of tech, legal, academia and policy reflect on what the EU AI Act has got right – and what could be improved.

The European business view

Kirill Skrygan
CEO of JetBrains

European businesses face far more regulation than Silicon Valley's startups, which makes it very difficult for them to compete. The compliance burden for firms in the EU ultimately slows the speed of tech development across the bloc.

The act is critical to ensure AI is transparent, controllable and, most importantly, not harmful. But we have to think about it together. The

speed of development of the AI industry in Europe, should be taken into account. We lack the intense collaboration between industry and government that would make the legislation more fruitful.

The EU should create exceptions, not based on employee headcount or revenue, but on [the substance of] AI initiatives. There are many tech products and services that

Europeans would like to develop inside huge corporations, and it's important to be able to compete with American companies and to preserve development velocity.

If AI startups aren't affecting a lot of people, we should leave them alone and let them grow. When the startup is significant enough to impact society, then let's start thinking about regulation.

The cybersecurity view

Dirk Schrader
CISO of EMEA at Netwrix

One of the most significant anticipated successes of the act is the standardisation of AI security across the EU. A key strength of the proposed regulation is its emphasis on security by design, mandating a lifecycle approach that integrates security considerations from the outset and throughout an AI system's operational life.

Several caveats could hinder the effectiveness of such AI-security

regulations. Threats are evolving rapidly. New attack vectors may emerge faster than static rules can be updated, meaning regular revisions will be required.

Enforcement will be difficult, too, as AI supply chains and cloud deployments cross jurisdictions. There's also the risk of 'compliance theatre', where firms prioritise tick-box compliance over meaningful security enhancements.

However, the act is the first major legislation to call out protections against data poisoning, model poisoning, adversarial examples, confidentiality attacks and model flaws. The real compliance burden will be determined by technical specifications that don't yet exist. These will define the practical meaning of "appropriate level of cybersecurity" and may evolve rapidly as AI threats mature.

The ethical view

Carissa Veliz, Professor of philosophy and ethics, University of Oxford



The EU AI Act is a step in the right direction. Every new technology must be accompanied by rules governing its use. The act is the best effort we have thus far.

I worry that an approach based on risk management is at odds with one that puts fundamental rights at the centre of liberal democracies. But the devil is in the details. The outcome will partly depend on how the law is interpreted, implemented and enforced.

The challenge is to stand out from ground as Europeans in the face of a complex geopolitical context, while also looking for international allies in other liberal democracies.

Everyone has a part to play in building the kind of society we want to live in – one that respects fundamental rights and freedoms. We need innovators to build competitive technology that's compatible with that. We deserve tech that is more respectful of our autonomy.

The policy view

Matthew McDermott
Director, Access Partnership



The EU AI Act is the first attempt to regulate AI comprehensively and that should be noted as a major achievement of the EU.

It sets the tone for global AI governance in the same way that GDPR set the tone for data protection. There's potential with the EU AI Act that other countries will follow that, too – it could become the blueprint. No other government will try anything so ambitious.

But, being the first act of its kind, it maybe gets things a little wrong on practicalities, particularly for businesses that work globally, for AI ecosystems that we haven't considered yet and also for SMEs.

What does it get right? The act is very clear on high-risk categories. Whether you agree or disagree with the list, its terms are easy to comprehend. It also gets its risk-based approach right for AI models, specifying unacceptable risk, high risk, limited risk and minimal risk. Moreover, the drafters have shown a willingness to listen to a broad range of stakeholders.

The compliance obligations are quite heavy, however, particularly for startups and SMEs that are building foundational models or sector-specific AI solutions. The reporting requirements are complex and, at times, ambiguous.

The legal view

Melissa Hall
Legal director at Morton Fraser



The EU AI Act provides clarity where legal obligations have been ambiguous. By establishing risk-based tiers, particularly for high-risk systems, it a framework in which to develop AI systems.

While the law technically gives open-source AI a pass, the most widely used systems will still fall under the regulation. This could help to boost public trust in AI, since popular open tools will still

need to meet safety standards. But it also creates additional costs for developers, who must ensure their tools comply with the regulation.

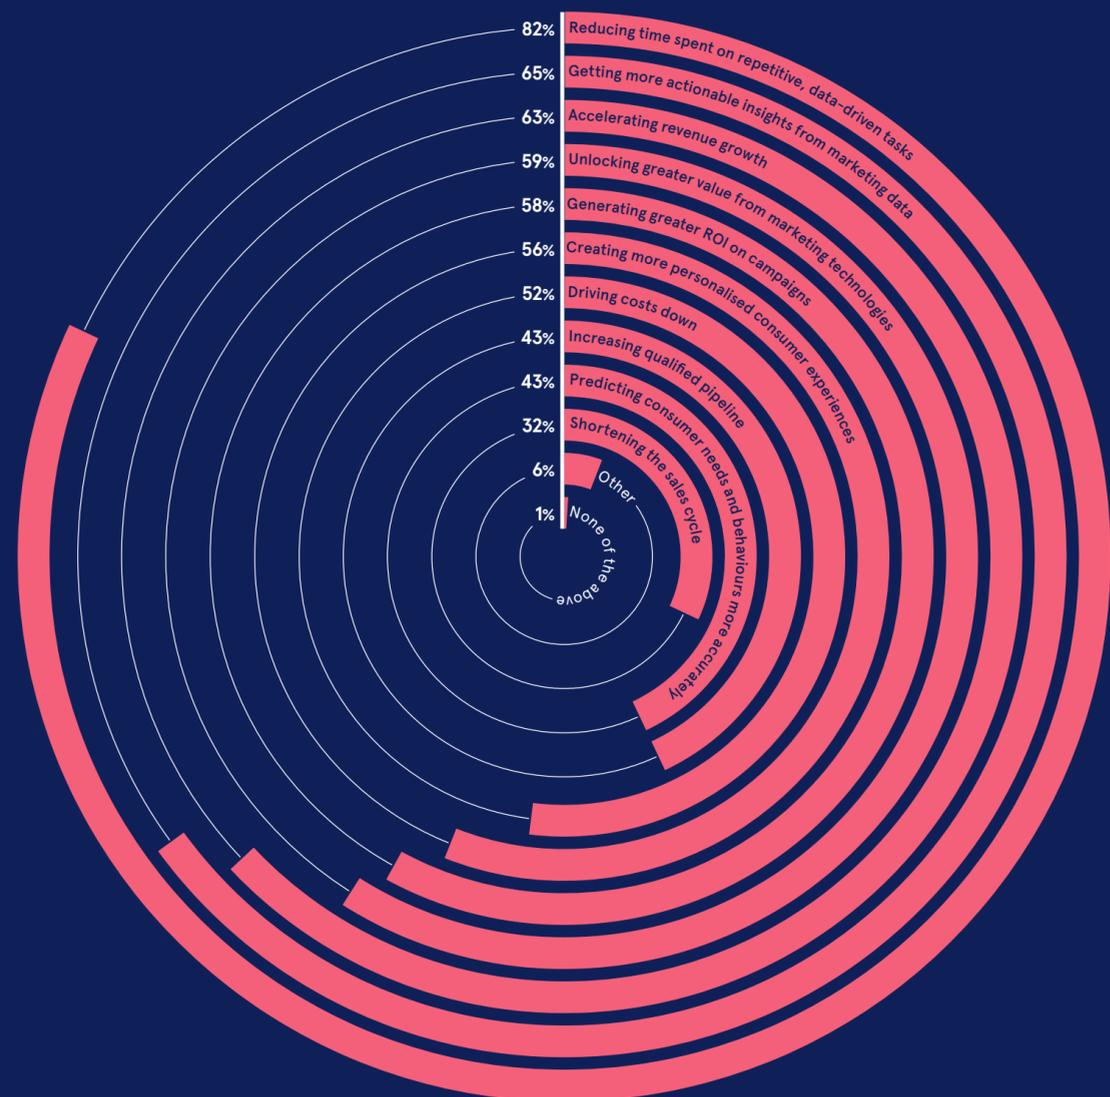
Ideally, the regulation would be clear and agile – protecting users and building trust, without creating a compliance burden that stifles innovation. The real test will be whether the EU can strike that balance and whether others choose to align or go their own way. ●

CMOS STRUGGLE WITH AI READINESS

Three-quarters of marketing leaders believe AI will be central to their team's success in 2026. But many are struggling to implement the technology effectively. There are plenty of reasons for this, but inadequate workforce training tops the list. It's a catch-22 for businesses seeking to deploy AI to drive efficiency and reduce costs. Employees must be trained to use the technology if adoption efforts are to be successful – but how many of those workers will be around much longer anyway? After all, over half of CMOs expect AI to eliminate jobs in their function.

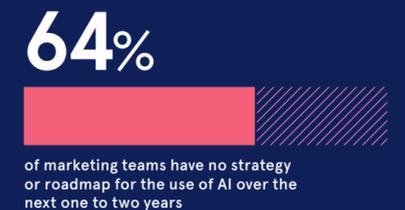
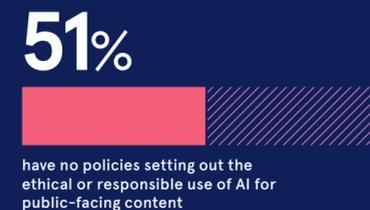
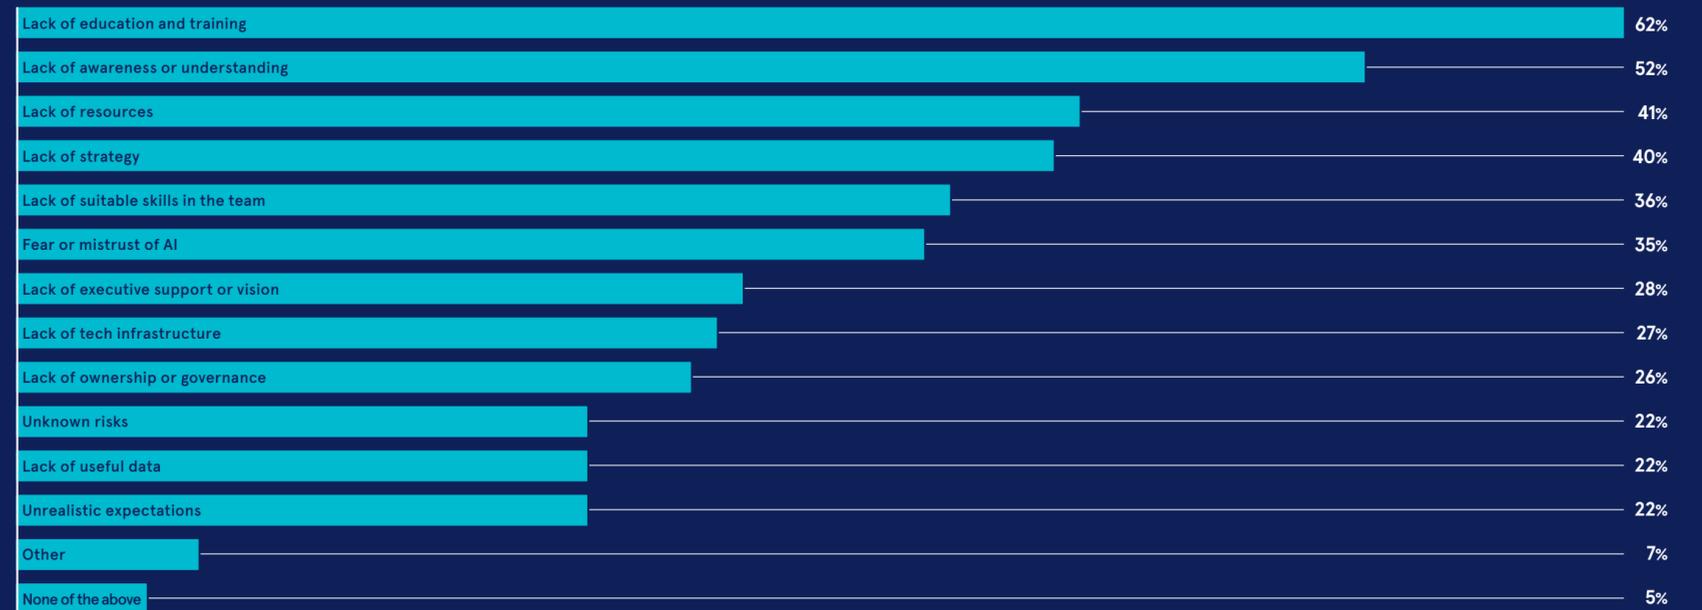
FOUR IN FIVE MARKETERS HOPE AI WILL REDUCE TIME SPENT ON REPETITIVE TASKS

Outcomes marketers are seeking to achieve with AI tools



A LACK OF EDUCATION AND TRAINING IS THE GREATEST BARRIER TO AI ADOPTION

Barriers to adopting AI tools in marketing



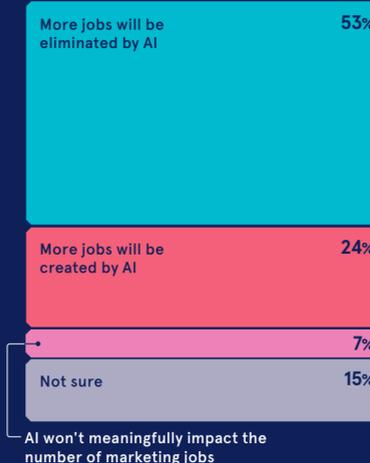
AT LEAST 65% OF FIRMS DO NOT OFFER AI-FOCUSED TRAINING

Organisations' provision of AI support to marketing teams



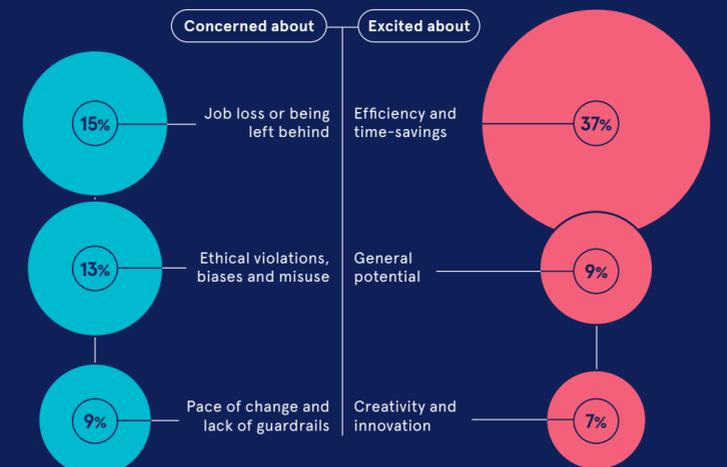
A MAJORITY OF MARKETERS BELIEVE AI WILL ELIMINATE JOBS IN THEIR INDUSTRY

Marketers' expectations of AI's impact on jobs in the industry



JOB LOSSES ARE NOT THEIR ONLY CONCERN

Marketers' top reason to be concerned or excited by the use of AI in the function



TRENDS

2025's biggest moments in AI

The whole world has gone mad for artificial intelligence. Here are the key moments, movers and shakers of AI in 2025 thus far

Tamlin Magee

A cursory glance at this year's headlines suggests that "technology" might as well be synonymous with "artificial intelligence". Even as many struggle to prove its value, business leaders can't seem to get enough of generative AI, eager as they are to alchemise their firms' pilot projects into something revolutionary.

Spending on the technology is through the roof. AI has rocketed up the agenda for governments, which increasingly view the tech as a matter of national security. The UK and the US, for instance, are developing ambitious sovereign-compute projects. Meanwhile, Silicon Valley firms are spending record amounts on AI infrastructure.

But the technology has not come without controversy. While more people than ever before are using AI tools, citizens are increasingly concerned about AI's impact on jobs and society as a whole.

The race for AI skills

With AI supposedly coming for everyone's jobs, from coders to poets, organisations are increasingly waking up to the need to upskill, well, everyone. For governments and enterprises of all sizes, preparing workers to use AI has become a priority in recent years.

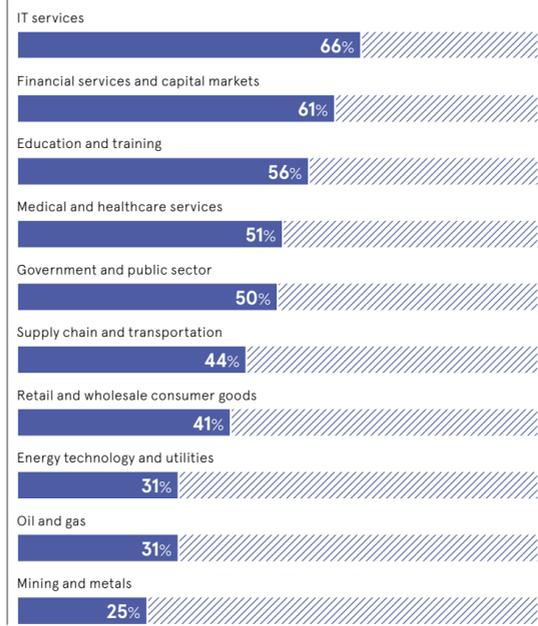
For instance, Infosys, an IT services and consulting firm, has introduced a layered upskilling programme, where employees receive training based on the level of AI expertise required for their job function. The programme classifies staff as AI "users", "builders" or "masters", and outlines different training modules for each group.

Meanwhile, in the Isle of Man, legislators are working to make AI training available to every citizen and business in the crown territory. And, across the UK, schools, colleges, community hubs and workplaces will soon offer AI-skills courses as part of a government-led initiative to train 7.5 million UK workers (about one-fifth of the workforce) to use GenAI effectively. Keir Starmer, the prime minister, has put AI at the heart of the government's industrial strategy.

THE SKILLS TO PAY THE BILLS

World Economic Forum, 2025

Share of global employers that view AI and big data as core skills for the next five years, by industry



The UK's AI action plan

When the Labour Party came to power in 2024, the new government recruited Matt Clifford, an investor and AI evangelist, to develop a blueprint for making the UK an AI superpower. Thus began the creation of the 'AI action plan', which was unveiled by Downing Street at the start of the year.

The plan will "mainline AI into the veins" of the UK, according to Starmer. Its flagship policies include loosening planning restrictions around so-called AI growth zones (clusters of data centres) and developing sovereign data networks and computing infrastructure.

Starmer also said that the UK must carve its own path with AI regulation, especially given that it now has the freedom to do so, post-Brexit.

Deepseek shakes World markets

When the Chinese AI startup Deepseek launched its R1 'reasoning' model, a system that is free to use and cheap to train, the international markets took notice. Tech leaders

had long assumed that LLMs require enormous amounts of energy and infrastructure to train and run, but the Deepseek model proved them wrong. Ripples were felt across the sector – the tech-heavy Nasdaq index plummeted by 3% – but worst affected was America's largest chipmaker, Nvidia, which shed nearly \$600bn (£484bn) in market capitalisation.

It wasn't long before the company recovered, however, as market observers noted that if AI can be

“Tech leaders had long assumed that LLMs require enormous amounts of energy and infrastructure to train and run, but the Deepseek model proved them wrong

made cheaper and more energy efficient, demand for it would likely increase in the long term. Providers would simply seek to accomplish more with the technology rather than use less of it. Nvidia is often considered a bellwether for the broader AI industry.

The Deepseek saga failed to undermine the business model of any major semiconductor firm, although shares in Broadcom and Taiwan Semiconductor Manufacturing Co also slid initially on the launch. It may, however, have ignited an AI arms race between the world's superpowers.

The market shock brought to light China's growing competency in developing AI models. Curious timing, given that Donald Trump, meanwhile, was preparing to launch a trade war intended at least in part to undermine competition against the US tech industry.

The regulatory picture

The EU became the first major trade bloc to legislate against potentially harmful uses of AI. The EU AI Act became law in 2024 and will be applied and enforced in stages. The

legislation aims to regulate transparency and accountability for AI systems and providers and sets out levels of acceptable risk for AI applications based on societal, ethical and legal considerations.

Like GDPR, the regulations apply to any organisation operating in the EU, so it's not only businesses located in the bloc that are impacted. EU member states will enforce the act and non-compliance carries hefty fines. The worst infringers will face penalties of €35m (£30m) or up to 7% of annual worldwide turnover, whichever is higher.

This year, the first and second stages of the act have come into force. The first establishes unacceptable, unethical uses of AI, including for mass surveillance or discrimination. The second targets general-purpose AI platforms such as ChatGPT, establishing a broad range of compliance requirements for companies that use LLMs.

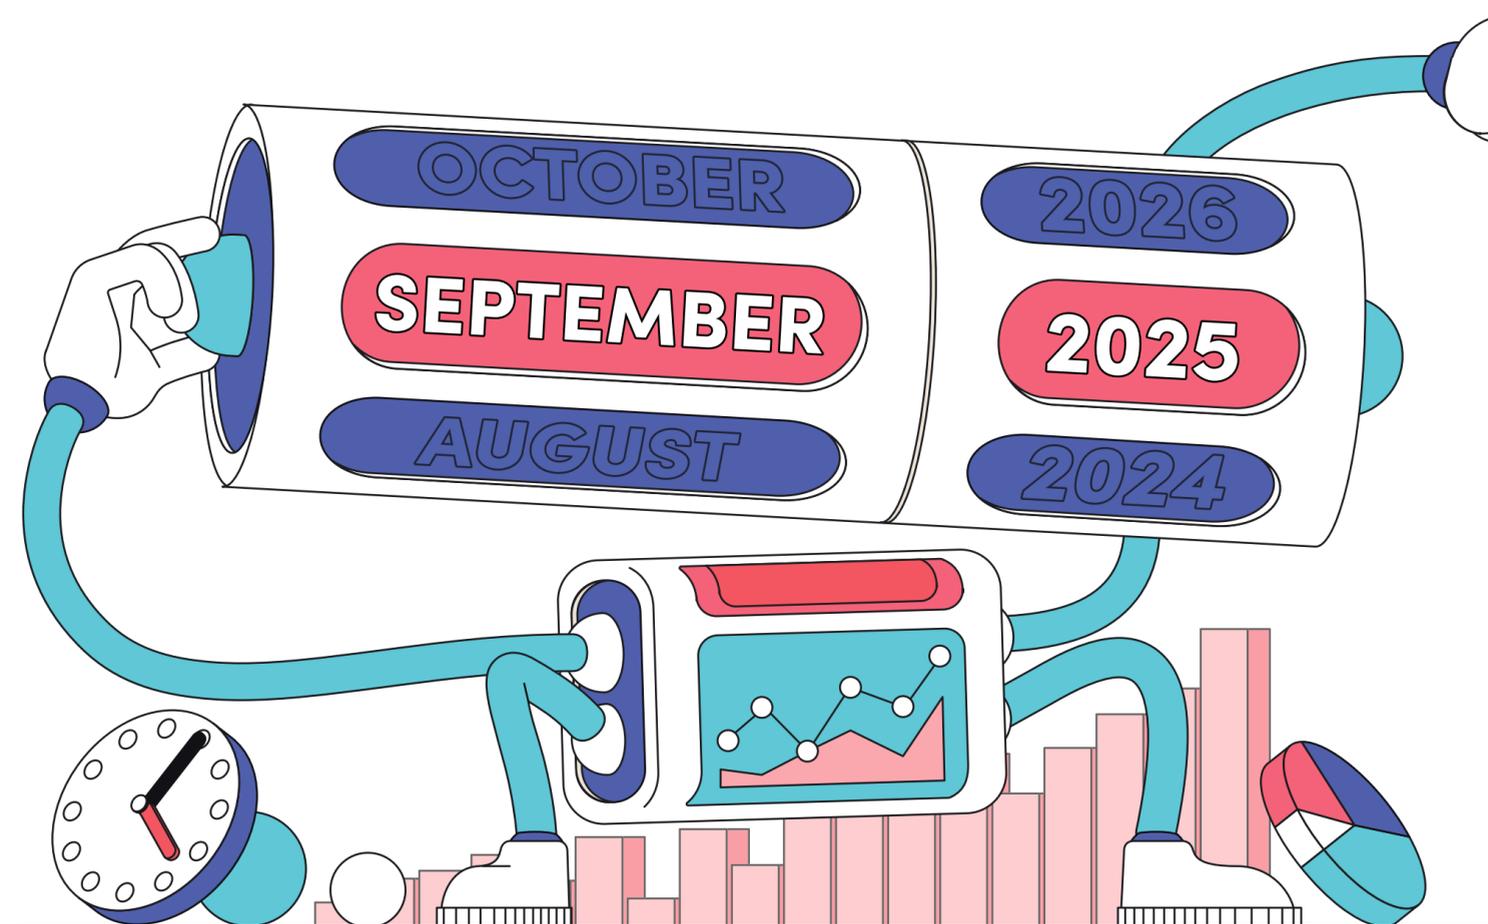
AI providers will be required to assess the safety of their models and demonstrate that they have done so through risk assessments and testing. They'll also have to retain technical documents on model architectures, make those records available to regulators and publicly disclose training data if

authorities deem it necessary. But some of ChatGPT's most loyal users were disappointed at GPT-5's lack of personality and generally unimpressive answers to prompts. OpenAI's chief executive, Sam Altman, claimed that a malfunctioning model-switching feature meant that users were unknowingly interacting with a "way dumber" version of the platform.

Competing models

While AI advocates excitedly await new models, some critics and sceptics have started to wonder how much further GenAI systems can be improved. If such tools have come as far as they can go, the AI bubble could be about to burst.

Still, OpenAI's popularity does not appear to be waning. The UK's former tech secretary, Peter Kyle, had apparently discussed a deal with Altman to provide a ChatGPT account to every UK citizen. ●



13 January UK AI action plan
Keir Starmer sets out Labour's vision to make the UK an "AI superpower", adopting a majority of policy items devised by Matt Clifford, the government's then-adviser on AI.

21 January Project Stargate
Donald Trump kicks off his second presidential term by announcing a \$500bn (£371bn) venture with OpenAI, Softbank and Oracle to develop AI infrastructure.

28 January Deepseek
Nvidia's stock tanks when word spreads about Deepseek, a Chinese GenAI startup that allegedly trained its LLMs using a fraction of the energy required by its Western rivals to train theirs.

2 February EU AI Act
The first provisions of the EU's sweeping AI Act come into force, banning dangerous applications of AI, such as for surveillance or social manipulation.

4 February Google ethics
Google removes from its AI principles a clause restricting its development of AI tech for weaponry or surveillance.

14 February Guardian + OpenAI
A report emerges that the Guardian encouraged strikebreakers to use GenAI to write headlines, following a journalist walkout in December over the publisher's sale of the Observer to Tortoise Media.

25 February Musicians' copyrights
Kate Bush, Damon Albarn and 1,000 other musicians put their names to a symbolic 'silent album' to protest the UK government's planned changes to copyright law, which would make AI's use of creative works opt out for artists, rather than opt in, by default.

4 March Oxford + OpenAI
Oxford University and OpenAI announce a five-year collaboration to provide students and faculty with access to "cutting-edge AI tools to enhance teaching, learning and research".

6 March Manus AI
The Chinese agentic-AI tool Manus AI becomes publicly available, amassing millions of pre-registrations and kicking off a wave of interest in autonomous AI agents.

17 March Tech-giant AI spending
A Bloomberg analysis finds that US hyperscalers are on track to spend \$371bn (£275) on computing resources and data centre infrastructure in 2025, a record high for annual AI investment.

25 March Gemini 2.5 launch
Google Deepmind releases its most powerful model yet, a so-called thinking model that can solve highly complex problems by "reasoning" through multiple decision paths before responding.

2 April Meta mass piracy
Writers attend a protest, arranged by The Society of Authors, outside Meta's London office over allegations that the firm pirated reams of copyrighted creative work to train its AI models.

20 May Google AI mode
Google searchers in the US are met with AI summaries at the top of their search results whether they like it or not. The feature would later become the default in India, the UK and 180 other countries.

20 June Softbank AI hub
Masayoshi Son, Softbank's chief executive, proposes a \$1tn (£742bn) AI-infrastructure hub to be developed in conjunction with Taiwan Semiconductor Manufacturing Co in the burning deserts of Arizona.

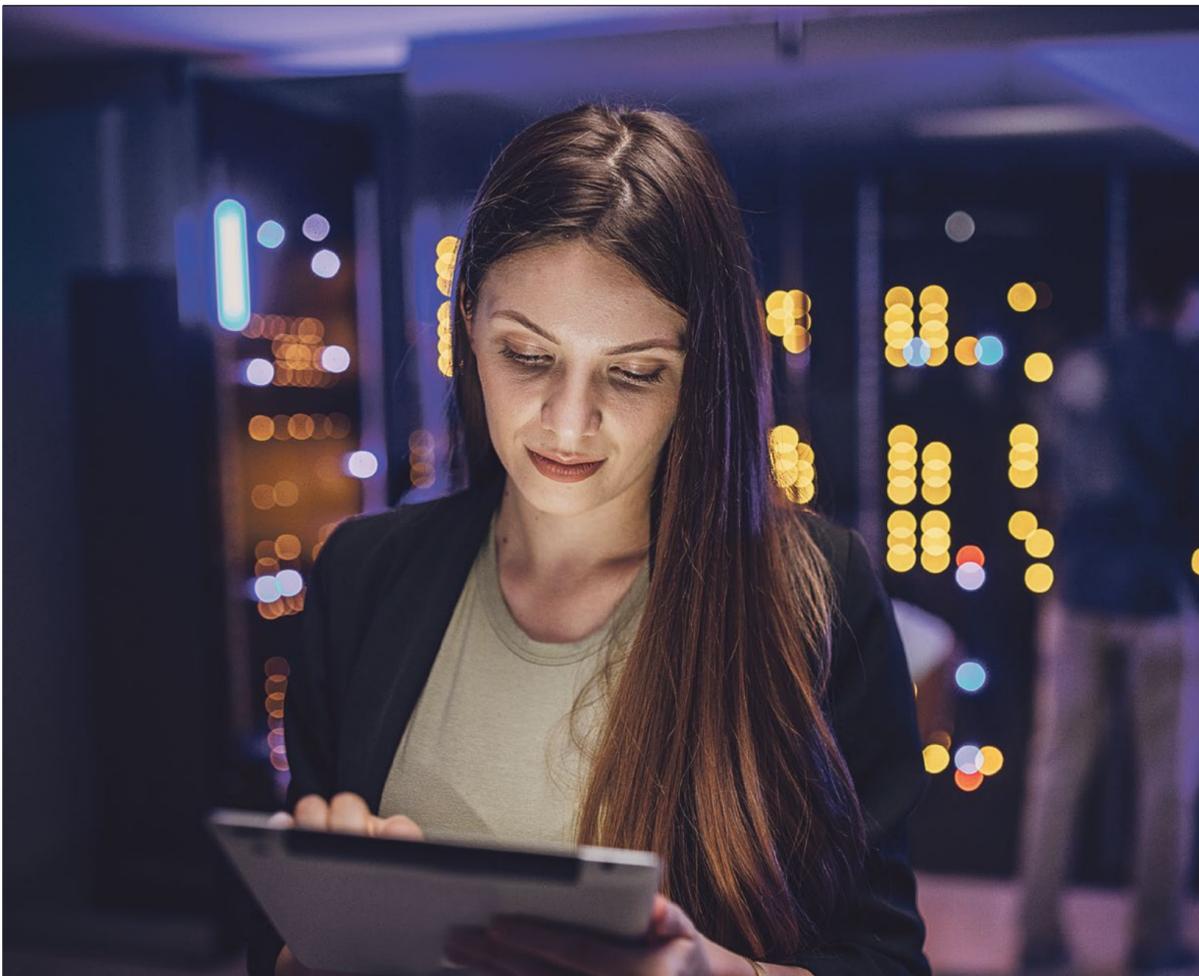
9 July Grok heils Hitler
Elon Musk comes under fire after Grok, the AI chatbot on his X platform, was caught praising the Nazi dictator Adolf Hitler, following an update to its system.

7 August GPT-5 launch
OpenAI hails the new GPT-5 as its most powerful model yet, but users say it's sterile and impersonal.

13 August Perplexity + Chrome
With US authorities calling for reforms at Google to address its alleged monopoly on online search, Perplexity, an AI-search platform, offers to buy the Chrome browser for \$34.5bn (£25.6bn). A US federal judge would later find no legal justification to force the sale of the browser.

16 August UK gov agentic
The UK's embrace of AI gets personal, as the government proposes deploying AI "helpers" for everyday citizens by 2027. The AI agents would assist people with everything from life admin to career choices.

17 September Microsoft boosts UK AI
Microsoft pledges \$30bn (£22bn) to AI infrastructure and operations in the UK, as part of a £31bn investment package with Nvidia and Nscale. The commitment comes off the back of Trump's second state visit to the UK.



Why enterprises need a flexible, scalable foundation for AI

As AI adoption accelerates, enterprises must build flexible, scalable foundations to unlock efficiency, resilience and long-term transformation

Interest and investment in AI are increasing as the technology becomes more accessible for a wide range of use cases. According to a survey by McKinsey & Company, most (92%) companies plan to accelerate their AI expenditure in the next three years. However, only 1% of leaders consider their organisations to be 'mature' in their deployment of the technology.

As businesses move beyond the experimentation phase in their digital transformation, they will increasingly look to 'enterprise AI' to develop business-specific workflows that can enhance employee efficiency and streamline processes.

The systems underpinning enterprise AI are fundamentally different from the non-proprietary, consumer solutions that many businesses have adopted following the boom of products such as ChatGPT, explains Vik Malyala. He is managing director and president

of EMEA and senior vice-president of technology and AI at Supermicro, a server-technology provider.

"Consumer AI serves generalised tasks at an individual level," Malyala says. "Enterprise AI is designed to address specific knowledge and workflows within an organisation. That can support everything from operations and customer service to supply chain management and real-time decision-making."

Enterprise AI infrastructures enable businesses to build agentic solutions in their workflows. Systems across every function of the business can act with significant autonomy, reducing the need for human intervention. While firms' AI strategies may differ considerably – some may train their own solutions, while others rely on fine-tuning existing models – the benefits of translating raw data into outcomes faster and more effectively can help organisations across all sectors.

"Organisations that advance from foundational models to building towards real-world deployment can reshape their day-to-day operations. Agentic AI can autonomously complete tasks, interact with users and adapt based on context. That promises to streamline operations, personalise customer experience or support predictive decision-making," says Ray Pang, senior vice-president of technology and business enablement at Supermicro.



AI is becoming a core driver of enterprise transformation. It will be embedded into every facet of the enterprise

"AI is becoming a core driver of enterprise transformation. It will be embedded into every facet of the enterprise. Almost every enterprise-level organisation is turning into an AI company. That transformation will be key in achieving long-term strategic goals," adds Malyala.

These outcomes are achievable with non-proprietary models, but building agentic AI capabilities in this way carries risk. Safeguards for data security and privacy are often inadequate and customisation limited because the technology has been built for broad uses. Governance is also unpredictable, as open models require frequent updates.

As such, the demand for localised, on-premise infrastructure accelerated by graphics processing units (GPU) is growing, particularly in industries where privacy, governance and compliance are critical. Organisations are looking for hardware solutions that can make integration scalable and accessible, reducing complexity and barriers to deployment while keeping costs as low as possible.

Scaling out enterprise AI infrastructure

As they iterate their solutions, enterprises also need adaptable hardware. Core to this technology are GPUs that use the peripheral component interconnect express (PCIe) interface, a high-speed standard in servers that enables the high-bandwidth, low-latency transfer of data between central processing units (CPUs), GPUs, memory and storage. The standardised nature of PCIe ensures compatibility and flexibility, allowing IT teams to optimise performance across a range of workloads. For scaling out beyond a single server, technologies such as InfiniBand and high-speed ethernet provide the external connectivity needed to link servers and data centres efficiently.

"Many businesses are exploring a modular approach," says Pang. "PCIe GPUs support enterprise workloads without the need for complex, highly-specialised components that can be integrated into existing infrastructure. That enables organisations to adapt and expand their computing needs as they evolve. Future-proof AI infrastructure means enterprises can experiment today and expand tomorrow – without a disruptive overhaul."

The development of enterprise AI in existing IT infrastructure requires careful planning. Unlike foundational AI, which often demands purpose-built facilities owing to its power, space, thermal and networking requirements, enterprise AI can be integrated into established systems. It is not a standalone environment but an added layer that must align with current resources to ensure scalability and efficiency without disrupting operations.

Pang says building proprietary enterprise AI therefore requires a holistic approach. Beyond modular servers and storage systems that can grow as AI workloads mature, businesses must also consider factors such as latency, data-pipeline readiness and sustainability and efficiency of power and cooling.

What might an ideal product mix look like? For a mid-sized company, which Pang defines as one with roughly 1,000

employees and an annual recurring revenue of \$250m, a good starting point would be an Nvidia-certified GPU server. These systems are validated to work seamlessly with Nvidia AI enterprise software and Nvidia networking, ensuring a reliable, full-stack solution for accelerated AI adoption. For example, Supermicro's servers containing either the Nvidia RTX Pro 6000 Blackwell Server Edition or Nvidia B200 GPUs enable both model fine-tuning and inference – the application of pre-trained models to unseen data to generate predictions – at scale.

In addition to GPU servers, companies need AI-optimised storage solutions to enable retrieval-augmented generation (RAG), which powers real-time, context-aware responses. To meet this need, Supermicro works with leading storage providers such as Vast, Weka and Nutanix, ensuring high-performance data access. High-speed ethernet switches provide low-latency connections between nodes and inference endpoints, allowing for fast transmission of data between GPU servers and edge devices. Finally, this hardware should be paired with a software stack that includes AI orchestrators – systems that manage and coordinate multiple tools, models and data pipelines – along with advanced RAG capabilities.

"That provides a systematic, automated solution to redeploy and refine models that maintain accuracy targets while lowering resource demands," says Malyala. "It also helps ease the risk in maintenance, customisation and governance. RAG and agent-advanced features will provide additional security and accuracy."

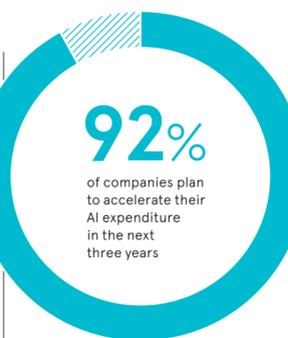
By combining the right infrastructure with an optimised software stack, enterprises can expand their capabilities, move closer to true agentic AI and rely less on generic, non-proprietary solutions.

Malyala adds: "It balances compute intensity, storage performance and energy efficiency. It can handle millions of customer interactions per month with reliability and compliance. The enterprise has the ability to own a proprietary customer service AI model, reducing dependence on generic, third-party SaaS providers."

There are some challenges in implementing these systems, however. For example, many on-premise data centres are only designed to deliver a maximum output of 20kW per rack, whereas racks in AI data centres can require up to 200kW of power each. Given the increased computing output of GPUs, that capacity can be quickly absorbed by fully-populated systems.

Cooling can also be a problem. Most enterprise data centres are air-cooled, which limits the types and density of GPUs that can be deployed. Plus, thanks to space constraints, it's often not possible to make major modifications. Enterprises must therefore adopt AI solutions that integrate into existing infrastructures.

Supermicro offers GPU-accelerated systems in a range of air-cooled form factors – including 5U, 4U and 2U



McKinsey & Company, 2025

– that fit into standard data centre racks. For organisations planning larger, purpose-built infrastructures, liquid cooling is an emerging option that can reduce operational costs and increase efficiency, though it typically requires new rack designs and dedicated facilities to support high-power CPUs and GPUs.

As enterprises look to mature in their deployment of AI, they will increasingly need to rely on the counsel of hardware providers that can deliver infrastructure capable of accelerating time-to-result and time-to-revenue for their proprietary agentic AI. Supermicro works closely with partners such as Nvidia to bring new platforms to market in weeks rather than months, adapting quickly to emerging hardware and unique customer needs.

"Through our global manufacturing presence and first-in-market support and system validation for new Nvidia GPUs, Supermicro helps organisations deploy enterprise AI infrastructure faster, speeding up our clients' adoption of AI and bringing faster return on investment," says Malyala.

For more information about how to scale the adoption of enterprise AI, visit [nvidia.com](https://www.supermicro.com)



Enterprise AI in action

From finance to telecoms, the uses for enterprise AI are rapidly expanding

Enterprise AI is moving from theory to practice, delivering measurable outcomes across industries. By embedding agentic AI into core operations, organisations can automate complex workflows, improve efficiency and unlock new opportunities for growth. Below are some of the most valuable uses for enterprise AI across key sectors.

Financial services

Banks and financial institutions are required to carry out extensive know-your-customer (KYC) and anti-money laundering activities. Yet these functions achieve staggeringly poor returns on their investments. By automating client onboarding, which triggers frequent KYC checks, and tracking irregular activities through agentic AI, firms can significantly reduce fraudulent activity.

Agentic AI can also streamline credit assessments for new customers.

Retail

Retailers have long strived for a true omnichannel experience. For customers, that is underpinned by the synchronisation of data across in-store and digital platforms. Agentic AI can support that aim by giving retailers real-time inventory visibility across stores, unifying pricing strategy and delivering consistent messaging on promotions on multiple channels. The ability to harness customer data can also drive marketing personalisation and higher sales, while reducing costs.

In-store automation is also helping retailers reduce labour costs, while improving customer experiences

Previously, such checks would rely on static data. But by using contemporaneous transaction data and economic indicators, lenders can continuously assess the risk on their balance sheets, creating dynamic lending models that can adjust in real time.

Manufacturing

According to research by IDS-INDATA, UK and European manufacturers are projected to lose more than £80bn to downtime in 2025. By collecting data from equipment sensors and production lines to identify patterns, agentic AI tools can detect early signs of wear and tear, predicting impending failures or disruption to production lines.

Disruption to supply chains also presents a significant risk. Agentic AI can predict supplier delays by tracking geopolitical or climate disruption, negotiate contracts with alternate vendors if there is a shortage of goods and balance production across factories to fulfill regional demand.

Intelligent models can also help manufacturers create more optimised designs of their final products, reducing operational and manufacturing costs, which in turn improves margins.

Agentic enterprise AI

The ability to draw from previously disparate data centres at speed will be crucial to improving customer experience, as businesses look to reduce cost and deliver highly personalised support and campaigns.

Agentic AI can also help to streamline internal processes in universal business operations. For example, in human resources, AI-powered technology can help to automate and personalise onboarding processes, deliver learning and development programmes, and improve collaboration across teams.



The ability to draw from previously disparate data centres at speed will be crucial to improving customer experience



OPINION

‘Outsourcing my writing meant outsourcing my thinking’

Writers may worry that AI is coming for their jobs. But those who have attempted to train the tech for written work know better, according to **Mike Reed**, co-founder of Reed Words



Mike Reed

Despite the astonishing pace at which LLMs are improving, many people say they can spot AI writing from its many ‘tells’, such as the gratuitous use of em-dashes. If only it were so simple. Not long ago, I spoke at the Scandinavian Creative AI Summit. Its founder, Lars Bastholm, had set me a challenge: “See if you can train an AI to write in your own voice and come speak about it.”

So I did – and it was enlightening. I used AI to write my LinkedIn posts for months, wondering if anyone would pull me up on it. They didn’t. Even seasoned writers who know me well didn’t spot it.

I started my experiment by getting ChatGPT, Claude and Gemini to pitch for the project. Each had to tell me why it was the one for the job – as well as which platform it thought objectively would perform best.

I picked Claude. It seemed like the best option for AI novices and the strongest on the nuances of language. I then asked Claude how I should train it for the task and followed its instructions.

It told me to first upload a selection of my work across various categories: articles, social posts and emails. Next, I had to identify patterns across those samples, including devices, style, tone and common phrases. Lastly, it asked me to create and upload a style guide.

With the groundwork complete, I was ready to let Claude loose. There were some simple rules: I would use AI to draft personal posts only, never client copy, and I would focus on LinkedIn. I post quite often there and many of my connections are copywriters – an especially AI-wary audience. Would any of them spot what I was doing?

My first Claude post took a lot of editing. It would have been much quicker to write it myself. After some extensive tinkering and prompt engineering I showed the copy to three creative directors at Reed Words, my writing agency, without mentioning the use of AI.

They liked it and had only minor suggestions. None saw any issues

“**Real writing isn’t about transcribing fully formed ideas. It’s a living process through which you generate, explore and refine those ideas**”

with tone or asked if I’d used AI. In fact, when I revealed what I was up to, they pointed out some phrases they said must have come from me. They were wrong.

After that, I used Claude to draft all my posts.

To be clear, this was a collaboration: I gave the AI plenty of feedback and rewrote some sections to train it further. But I tried to keep as much of Claude’s copy as possible. And I kept hearing the same thing about my undisclosed use of AI: absolutely nothing.

The first piece I published virtually unchanged was a story about AI. It didn’t just work – it out-performed all my previous posts by a wide margin, with three times as many impressions as its nearest rival and twice as many reactions.

That may have been because the post was about AI, which is a hot topic in every field. And I suspect Claude also incorporated an understanding of what works best on a platform such as LinkedIn. A colleague observed that the post sounded like a “LinkedInified” version of me. But based on engagement, at least, AI-Mike was doing a better job than human-Mike.

Things got trickier when I went long-form. I tried to write an article for LinkedIn based on a recent conversation with a new client. I threw some bullet points at Claude and it spat out a draft. I wasn’t very happy, so I tried again. And again. But the piece remained rambling and vague. It just wasn’t capturing what I wanted to say.

Pondering this, I realised Claude wasn’t the problem – I was. I didn’t know what I wanted to say or how to say it, because I usually work that stuff out by writing. The novelist EM Forster once asked: “How do I know what I think until I hear what I say?” Real writing isn’t about transcribing fully formed ideas. It’s a living process through which you gener-

“**Whatever it’s doing might look like writing, but it’s in fact something else**”

ate, explore and refine those ideas. It’s a way to find new ideas and connections, explore counterpoints and spot gaps.

The critic Geoff Dyer said that writing “encourages me to write things I hadn’t previously noticed, to have thoughts I hadn’t had before writing began”.

Outsourcing my writing to Claude robbed me of this vital process.

AI can analyse, synthesise, summarise – even ventriloquise. But it can’t have new ideas. Claude might be able to summarise what Forster meant, but it cannot come up with the original thought.

The problem is that AI drafts only appear to be polished. It’s slick writing, quickly generated. But this is a trap. There’s little substance beneath the surface, because you’ve killed the thinking process before it got started.

In reality, an AI can only work with what you feed it. It can neatly package what’s there but it can’t develop it or suddenly realise that everything after the second paragraph is crap. Whatever it’s doing might look like writing, but it’s in fact something else. Beneath that surface sheen, the content is often sketchy, thin and unoriginal.

Outsourcing my writing meant outsourcing my thinking. And that meant never getting past the germ of an idea. Claude’s neat drafts became a sort of cage, with no space to pursue, test or expand those ideas.

This superficiality is a problem for brands that think they can simply train an AI in their tone of voice and get it to deliver content without spending time or money on people like me.

But the danger is that they end up endlessly regurgitating the same old stuff. People will get bored of that, even if they don’t recognise the writing as AI. Boredom is bad for any brand, so at some point they’ll have to address it.

To answer Lars’s question: yes. I can absolutely train an AI to write in my voice. I’ve shown that I can nudge it into generating virtually instant content that few readers, if any, could spot wasn’t mine.

But this is mere mimicry. Real writing is about pulling on a thread to see where it leads. AI snips that thread. If you accept what it spits out, you’ll never know how many richer, more original ideas there were, further along that thread.

It is richness and originality that sets people apart from AI. If you give those up, what’s to stop you from being replaced?

I will, of course, keep experimenting with AI. It’s very useful and only getting better. But I’ll also keep writing and thinking for myself. And I believe it’s important we all do. ●

Co-innovation in action: solving AI’s challenges with open collaboration

By orchestrating open collaboration across hardware, software, cloud and integrator partners, Red Hat is helping businesses overcome AI barriers

AI is evolving at breakneck speed. In just a few years, it has transformed how businesses make decisions, personalise customer experiences and design new services. However, while some firms have successfully integrated AI across their workforce and workflows, others are struggling to transition from proof-of-concept experiments to production deployments.

Three main issues hold them back. Firstly, there’s the high cost of inference – the process of using a trained model to generate outputs – when running large, proprietary GenAI models and model services in production, with costs escalating as organisations scale and bring more use cases online. Secondly, there’s the complexity involved in selecting and aligning these models with organisational data, either through retrieval augmented generation (RAG), which involves combining them with external knowledge sources to provide more accurate results, or expensive fine-tuning. And lastly, there’s the challenge of deploying and managing AI across hybrid cloud environments.

“It can be difficult to deploy AI where you need to – but at the same time businesses need to embrace that kind of flexibility,” says Martin Isaksson, go-to-market lead in the AI business unit at Red Hat, an enterprise open-source software provider.

The speed of innovation today and complexity involved in deploying AI at scale are challenges that are best addressed collaboratively. Businesses must harness the unique expertise of a broad ecosystem of partners, from hardware and software vendors to cloud hyperscalers and system integrators, to deploy AI effectively across all the environments where it’s needed. And this requires a consistent, open platform for co-innovation.

89%
of organisations that have adopted AI use open-source AI in some form in their infrastructure

60%
of business decision-makers report significant cost savings associated with open-source AI
Linux Foundation, 2025

“That sort of foundation is incredibly important, because no single vendor can provide every tool that’s needed,” says Isaksson. “A key benefit of working with open source is that you’re always operating at the same pace as innovation, right where it happens.”

This approach underpins Red Hat’s AI portfolio. It includes Red Hat Enterprise Linux (RHEL), a foundation model platform for developing, testing and deploying LLMs, and Red Hat OpenShift AI, a platform for managing the entire lifecycle of AI and machine learning models, from development and training to deployment and monitoring, across hybrid cloud environments.

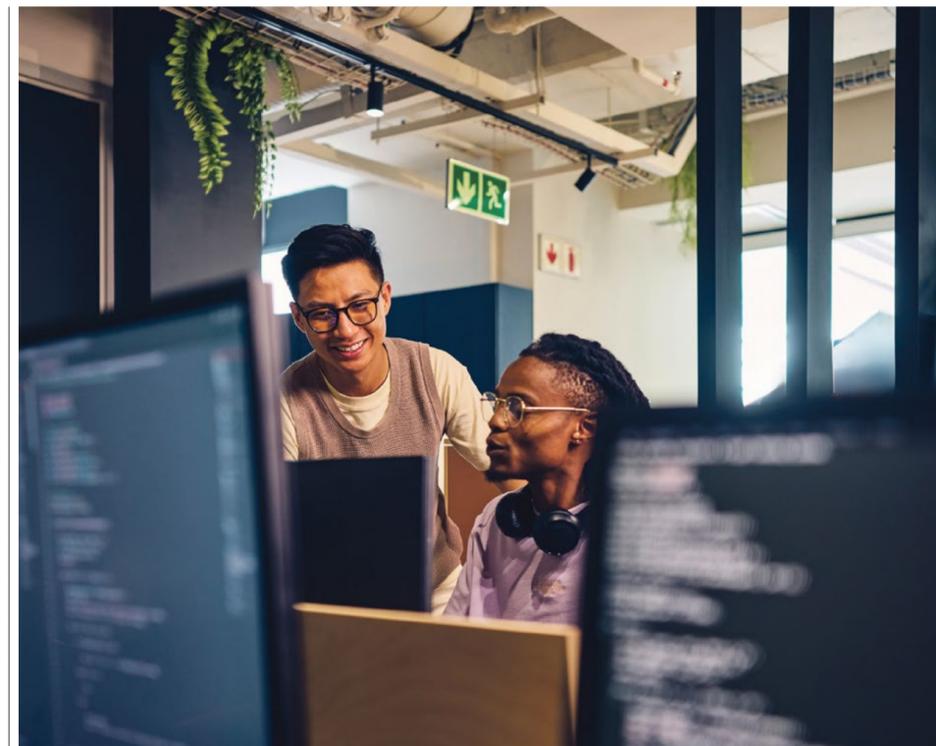
Red Hat’s partner ecosystem also includes global system integrators such as Accenture and Wipro, hardware providers including AMD and Nvidia, and cloud specialists such as IBM and Google. By orchestrating value across them, Red Hat enables businesses to accelerate AI initiatives without becoming locked into a single vendor.

Any model, any accelerator, any cloud

With so many LLMs, inference server settings and accelerator options available today, businesses need an easy way to navigate them and ensure that tradeoffs between performance, accuracy and cost meet their needs. “The ability to easily switch between different models is of great importance,” says Isaksson. “We have a pre-selected model catalogue with optimised models, so they’re faster and cheaper to run.”

These curated and validated models are available on Hugging Face, an open-source community for co-innovating models, datasets and applications. Deploying them helps businesses to reduce their dependence on proprietary LLM providers, whose solutions often include “black box” algorithms and training data. “We believe you should be in control of the infrastructure, the data, the model, how it works – and in the end the whole application,” says Isaksson.

Small language models (SLMs), a subset of LLMs that can quickly and easily be customised with enterprise data for specific tasks, are another key area of focus. Together with IBM, Red Hat co-created InstructLab, an open-source project designed to lower the barriers to customisation by enabling domain experts, not just data scientists, to fine-tune SLMs using their own knowledge and data. In fact, when paired with RAG, the right SLM could even outperform a proprietary LLM.



“**A key benefit of working with open source is that you’re always operating at the same pace as innovation, right where it happens**”

Faster, more efficient inference has also emerged as an important element of successful AI strategies. Red Hat AI Inference Server optimises model use of inference across the hybrid cloud to drive down costs. Built on the open-source LLM project, it can support any GenAI model, on any AI accelerator, in any cloud environment.

Scale and trust

Red Hat’s OpenShift AI has helped data scientists at DenizBank, a private bank based in Turkey, to reduce model development time, for instance, freeing up their data-science teams to create new business value instead of managing infrastructure. Integration with hardware-accelerator dashboards has helped to optimise GPU use, with the platform automatically scaling up the slices of GPU a model has access to, as needed. This enables

more workloads to run simultaneously without the need for additional GPU hardware, maximising the return on their existing hardware investments and improving efficiency.

Such AI solutions can help businesses to optimise the performance of AI deployments across a range of hardware configurations. “If your accelerator infrastructure is fragmented, with resources scattered in different places – some in the cloud, some on-prem – with one platform you can virtually pool all these resources and optimise them,” Isaksson explains.

Open-source co-innovation also enables businesses to stay on the right side of rapidly evolving AI governance and security requirements. Use of proprietary LLMs often raises concerns about security, privacy and safety, while uncertainty around the training data used and the accuracy of responses can increase legal risks for businesses. Open-source projects, on the other hand, provide the transparency needed to identify bias and privacy issues before they become problematic. This level of transparency is critical for managing the significant legal and reputational risks associated with enterprise AI.

“That’s really important,” says Isaksson. “This is why we’re supporting TrustyAI, which is an open-source responsible-AI toolkit that aims to solve AI’s well-documented problems with bias.”

The TrustyAI community maintains several responsible-AI projects, involving model explainability, model monitoring and responsible model serving. Red Hat engineers in the community recently developed safeguards that ensure LLMs behave ethically, safely and within organisational or regulatory boundaries, making them more viable for high-stakes deployments. The company also plays an active role in Llama Stack, Meta’s open-source framework for building GenAI applications, and supports Anthropic’s so-called Model Context Protocol, which standardises agent-to-application interactions.

Ultimately, businesses need innovative solutions that will help them overcome AI’s cost and complexity barriers. By providing the open, hybrid-cloud platform to unify and orchestrate this rich ecosystem of partners, businesses can turn experiments into scaled solutions, successfully deploy AI across hybrid cloud environments and retain control over their data and costs. In other words, they can ensure that their AI systems work on their terms.

For more information please visit redhat.com





From busywork to breakthrough: rethinking work in the AI era

To unlock the full benefits of AI, organisations must ensure it works at the team level

Innovation isn't just about ideas, it's about momentum. But it's hard to maintain the drive when working days are filled with administrative tasks, meetings and other interruptions. According to a recent report from Miro, for every hour spent on strategic or creative work, knowledge workers spend three hours on administrative tasks, meetings and communications. Six in 10 workers say these momentum blockers sap their concentration, reduce their creativity and raise their stress levels. Teamwork and collaboration suffer too: 70% of respondents experience collaboration issues at least monthly, and 55% said projects and initiatives at their company often lose momentum.

Excessive busywork is also tied to organisation-level problems, such as data and communication siloes. Constant emails and messages ping-pong across different channels and teams can create communication chaos. Ceaseless documentation and reporting tasks, and information that sprawls across different tools and

apps, also leads to duplicated work, missing data and context and slower decision-making.

AI promises to address many of these issues and create more time for the momentum work that drives innovation. It presents an opportunity for organisations to fundamentally rethink how work gets done, which could improve both the wellbeing of employees and teams' ability to innovate at speed. But it isn't a silver bullet. Poor implementation can compound many issues with existing working practices.

"It's not about taking a new technology and tacking it on to existing processes," says Grisha Pavlotsky, chief transformation officer at Miro. "It's an opportunity for leaders to ask: 'Is my operating model right for this new era? Are processes and roles still valid?'"

Improving teamwork

Major technology shifts of the past, such as electricity transforming how factories operated, highlight the need to rethink work rather than incrementally improve it. Yet many recent implementations of AI in the workplace have

Organisations that fail to ask these questions risk thinking they're moving faster while not actually progressing much at all. They may achieve high AI-adoption rates, for example, yet see only relatively small productivity improvements at the team and organisational level.

"You need to give AI to the team, not to individuals," says Pavlotsky. "If you miss that nuance, you will be fooled by the fact that AI adoption is very high in your organisation. Yet in reality you're not going to be identifying the right problems to solve, and you're not going to be experimenting fast enough."

Organisations often see AI as a means to speed up the delivery phase of product development. But Pavlotsky argues that many of the most transformative opportunities lie earlier in the development life cycle, during the discovery and definition phase.

Indeed, in a world where the cost of building software is plummeting thanks to AI, the ability to iterate rapidly on concepts, test assumptions quickly and reach high-quality specifications faster is becoming a key competitive battleground.

"We've figured out how to optimise the delivery side of product development," says Pavlotsky. "Tools like GitHub Copilot and Codex have solved that. So the key question now is not, 'Can I build something?' It's, 'Should I?'"

AI tools can synthesise data and identify the key insights, brainstorm ideas for improvements and even prototype a new user experience to generate feedback. These AI sidekicks – intelligent agents that work alongside teams – can even provide instant access to specialised expertise by emulating the response of a senior leader, helping teams move through the early stages of product development at greater speed. Ultimately, they should be able to iterate and test many more ideas, enabling them to reach product-market fit much faster.

However, none of this is possible without AI that has team-level shared context. "If everyone uses their own AI individually, then in meetings you will just argue about prompts and contexts. 'What did you feed the model? What version of the file did you use?' That's wasted time and the loudest voice may dominate, leading to misguided decisions."

While this shift will require new competencies – and leaders to have a clearer understanding of the value of certain artifacts – it could also help to free workers from the confines of their current roles, which are far too focused on maintenance work today.

The organisations that thrive in the AI era won't be those that simply deploy AI tools to help individuals carry out existing – and often draining and unfulfilling – tasks more efficiently. They'll be the ones that use game-changing technology to fundamentally rethink how work gets done.

The right context

Context engineering is an important new discipline that addresses this issue. It's about ensuring AI has the right shared, dynamic inputs. Because just as humans need shared context when they solve the problem, team-level AI also needs it to provide consistent, high-quality outcomes.

Team-level AI that draws input from a collaborative workspace can also serve every member of the team with outputs in their preferred form factors: technical diagrams, documentation, data tables, slides, prototypes and so on. But while the outputs may differ, all of them are based on that same shared context.

Crucially, everyone must be able to see and curate this evolving context in real-time. "Most current tools let you upload files that are static at that moment in time," says Pavlotsky. "But if you're working on Miro canvas, you have a lot



Miro, 2025

of context there that the AI is going to be using for a task. It's live, it's dynamic and it's visible, so everybody on the team understands what the AI is using. Furthermore, Miro AI can ingest context in all types of form factors: visual, docs, tables, dashboards and diagrams."

Smarter use of AI at the team level will also enable human workers to shift from specialists focused on narrow domains to orchestrators of end-to-end value creation. However, this requires an understanding of how different work streams connect, what intermediate outputs feed into final deliverables and how quality is maintained across complex workflows.

"You're not just responsible for one intermediate artifact – you're responsible for the final artifact and you need to understand what it takes to get there," says Pavlotsky.

While this shift will require new competencies – and leaders to have a clearer understanding of the value of certain artifacts – it could also help to free workers from the confines of their current roles, which are far too focused on maintenance work today.

The organisations that thrive in the AI era won't be those that simply deploy AI tools to help individuals carry out existing – and often draining and unfulfilling – tasks more efficiently. They'll be the ones that use game-changing technology to fundamentally rethink how work gets done.



Explore how organisations are using Miro to unlock the full potential of AI



FUNDING

Living on the hedge: investors exploit AI bubble

A 23-year-old tech specialist with no prior Wall Street experience has raised \$1.5bn (£1.1bn) for his AI-focused hedge fund. Will others follow his lead?

Sam Birchall

A 23-year-old with no professional investing experience has launched a \$1.5bn (£1.1bn) hedge fund in less than a year. That's faster than veteran portfolio managers can unleash their most prized investment vehicles. What's more, the AI-focused fund returned 47% in the first six months of 2025 – trouncing the S&P 500's 6% gain.

Soaring valuations are fuelling hedge funds betting on the AI boom. And some tech insiders may have an edge over seasoned finance professionals. So who is this AI analyst making waves on Wall Street?

Leopold Aschenbrenner is the founder of the San Francisco-based fund, Situational Awareness. He grew up in Germany and graduated from Columbia University as valedictorian at age 19. He then worked briefly as a researcher at OpenAI.

Aschenbrenner first announced his Wall Street ambitions in a provocative 165-page manifesto on the potential and risks of AI. "Everyone is now talking about AI but few have the faintest glimmer of what is about to hit them," he claims. "There are perhaps a few hundred people in the world who understand just how crazy things are about to get, who have the situational awareness."

Situational Awareness is betting heavily on infrastructure needed to power AI systems just as businesses prepare to pour millions into essential components, such as high-performance chips. It is targeting companies that will benefit from AI's build-out, while shorting industries likely to be left behind.

AI products, Aschenbrenner claims, will become the "biggest revenue driver for America's largest corporations, and by far their biggest area of growth". The first \$10tn company, he says, could emerge very soon and, when it does, big tech will funnel hundreds of billions of dollars into growing AI infrastructure, catalysing an "extraordinary techno-capital acceleration".

The pitch has resonated with some of the most influential figures in the tech sector. Meta's Daniel Gross, Patrick and John Collison, the founders of Stripe, and Nat Fried-

man, the former GitHub chief are backing Aschenbrenner.

And the buzz, as well as the early success, around Aschenbrenner's fund signals a major shift in the wider financial markets.

Situational Awareness is a pure-play bet on AI infrastructure and critical components, including chips, data centers and power generation. Investors typically treat these as support sectors within broader categories such as tech or industrials. But thematic funds are now pushing capital to a small number of managers who have articulated clear paths forward.

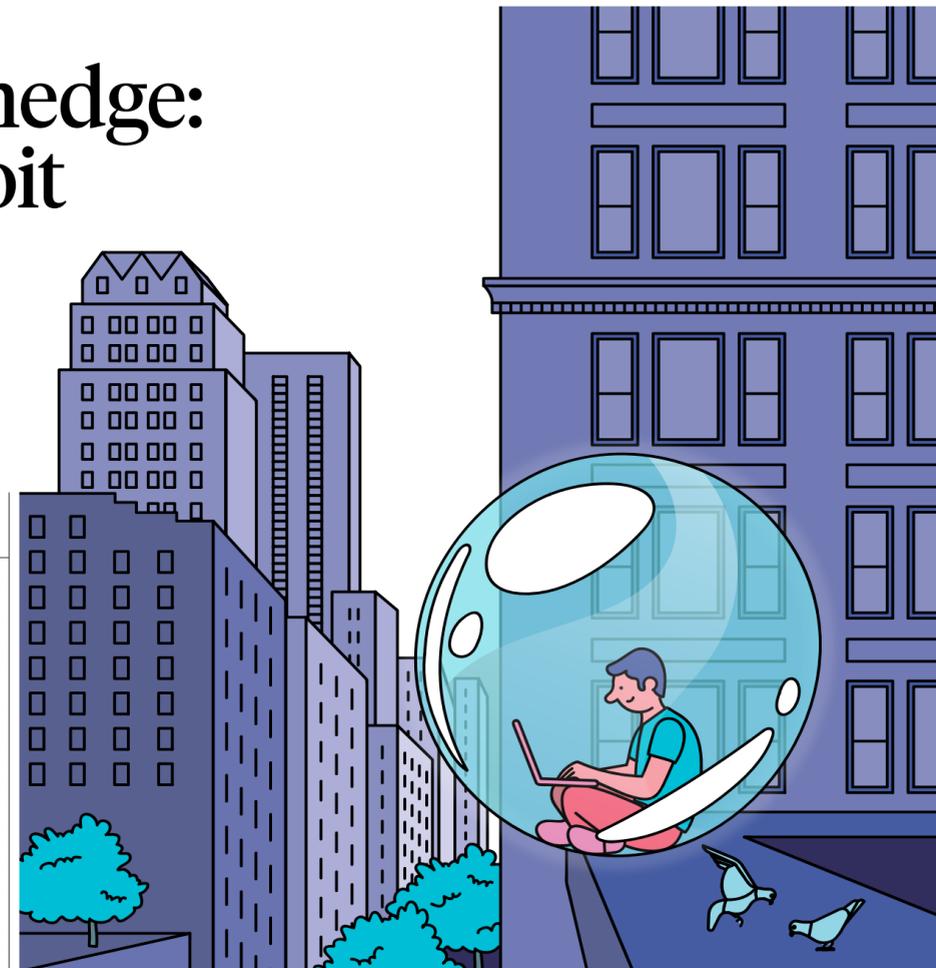
Other AI-focused hedge funds, such as Value Aligned Research Advisors and Point72 Asset Management, have also amassed billions in contributions, reflecting growing demand for AI bets despite market volatility and overlapping positions in AI-adjacent companies. Such funds may help to write the next chapter of AI investing. But investor sentiment can sour quickly. Extreme thematic booms often end with sharp corrections. Consider the dot-com stocks in the late 1990s or renewable energy in the early 2010s.

The Deepseek selloff earlier this year highlights the risks of over-concentration. The unstoppable momentum fuelled by a few high-flying stocks can quickly give way to a glaring vulnerability if those stocks fail.

But Aschenbrenner proves that experience isn't everything in today's investment market. An unconventional founder with no prior Wall Street background, his deep domain expertise has proved compelling for investors. What might once have been seen as a liability among investors could now become a critical edge for new fund managers in the AI space.

Behind the scenes, a shift in capital flows is amplifying the AI boom. As traditional banks retreat from financing certain types of projects, private credit funds have rushed in to fill the gap.

Private loans to the technology sector increased by \$100bn (£74bn) over the past 12 months, reaching a total of \$450bn (£334bn). Much of



“Everyone is now talking about AI but few have the faintest glimmer of what is about to hit them”

that has been directed toward the data centre build-outs of Meta, Amazon and OpenAI.

Non-bank financing is buying Silicon Valley's commitment to AI, but relying on it creates plenty of risks. Analysts say the boom is driven by opaque and untested borrowing, making the sector more fragile and harder to regulate.

The willingness of investors to prioritise vision over track record can accelerate growth, but it can also amplify risk.

OpenAI's Sam Altman has described AI as both a "bubble" and "the most important thing to happen in a very long time". Investors, he argues, are "chasing a kernel of truth" but pushing valuations into "insane" territory.

Economists have spotted some worrying signs, too. Across the sector, price-to-earnings ratios are stretched to levels reminiscent of the late-1990s, when enthusiasm for internet-based companies led to a stock market crash. JP Morgan estimates that spending on data centres now contributes to a disproportionately high share of US GDP growth – about 10 to 20 basis points. Still, Altman insists that trillions will be spent on data infrastructure "in the

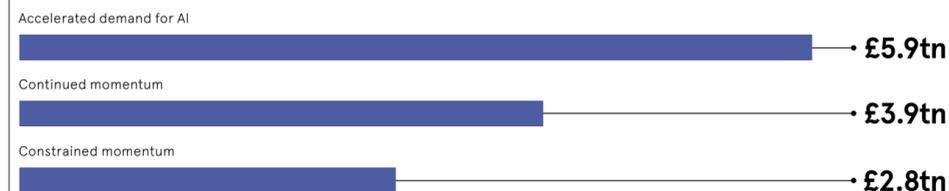
not very distant future" – whether economists and sceptics deem it reckless or not.

The economist John Maynard Keynes warned that "markets can remain irrational longer than you can remain solvent." Much like the tech boom in the 1990s, AI technology is generating massive amounts of capital into the sector. Even if the broader economy enters a slow-down or tips into recession, AI stocks will likely continue to climb, fuelled by narrative momentum and investor conviction.

But eventually reality will catch up. At some point, investors will demand proof that AI companies generate profits, not just attract capital. And when that day comes, AI providers will have to reconcile with good-old-fashioned finance fundamentals: revenue, margins and profitability. ●

NON-STOP FUNDING

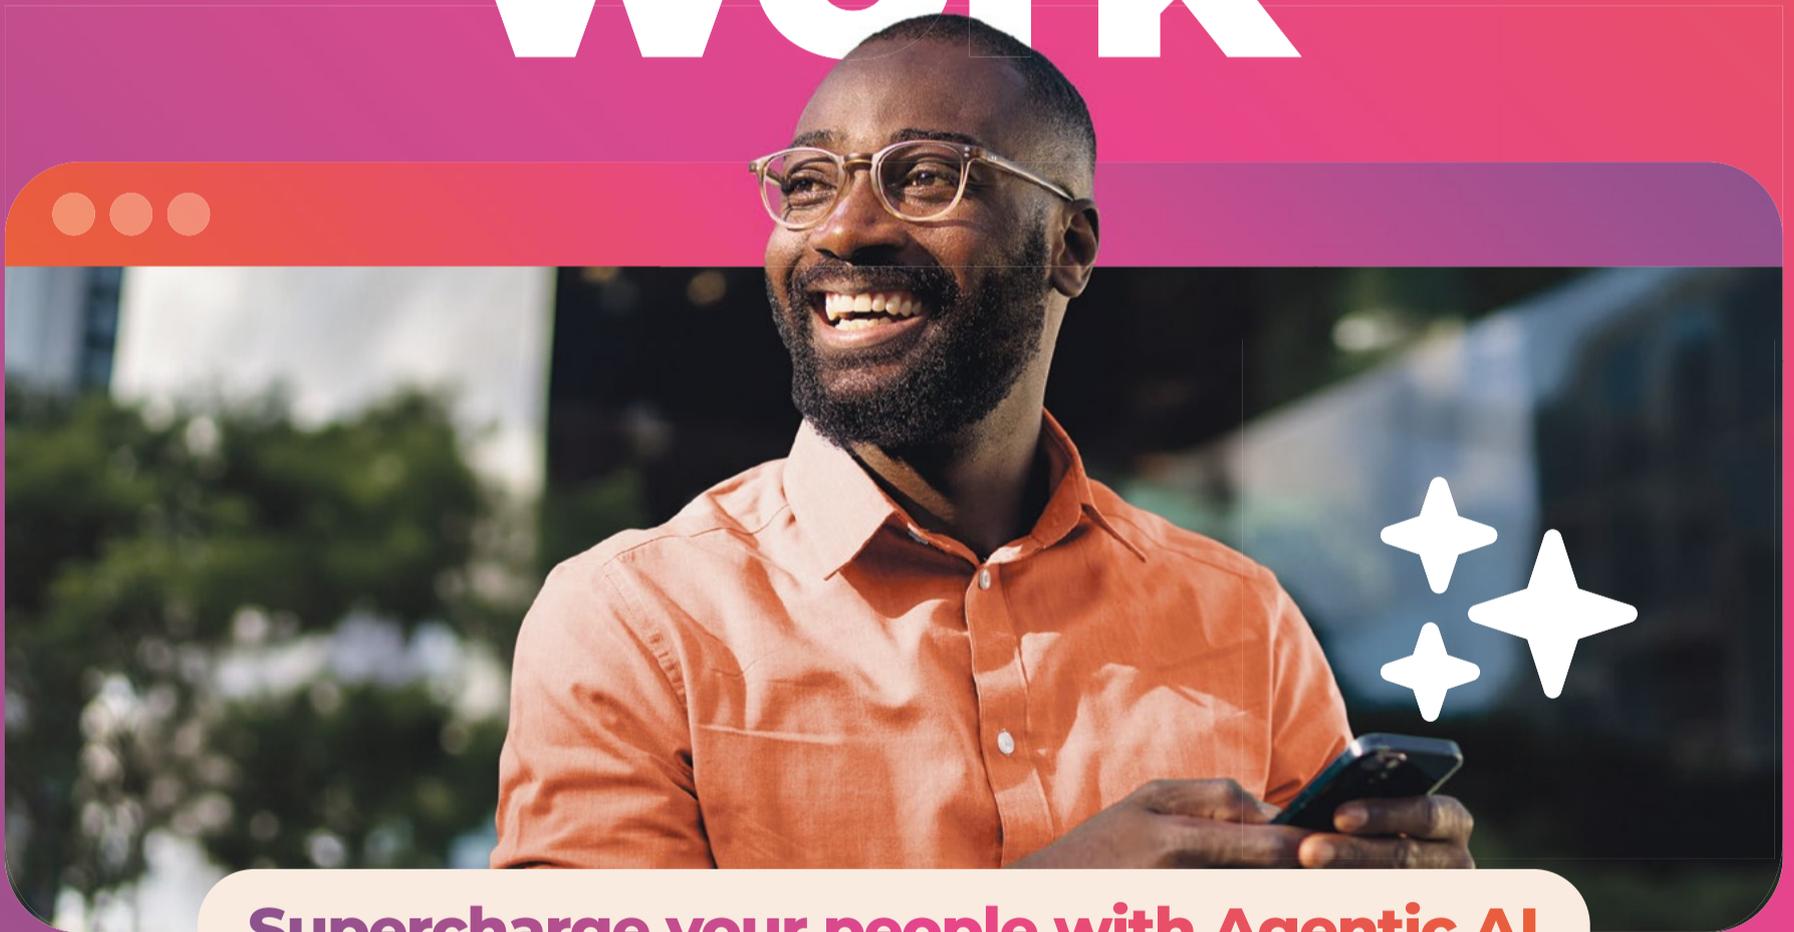
Projected worldwide spending on AI data centres, IT equipment and power from 2025-30, by scenario



McKinsey, 2025

oneAdvanced

It's AI in a day's work



Supercharge your people with Agentic AI

Less work woes. More work flows
with AI-driven reporting, streamlined
scheduling and automated data analysis.



Scan to
discover
more