

The C-suite Agenda


- p4** Executive risk in the age of autonomy
- p10** Deep tech's hard reality: why capital is shifting to infrastructure
- p2** Europe closes the door: soaring costs are forcing strategic sectors to go local



 | oneAdvanced

Make work flow

One trusted, intelligent system of work, connecting people, processes, data and AI.



oneadvanced.com

Contributors

Simon Chandler
Technology journalist whose credits include Wired, Digital Trends, Business Insider and Forbes.

Tom Dennis
Editor and journalist covering tech and innovation for almost two decades, from consumer gadgets to enterprise stacks.

Jay Lesada
Staff writer covering the enterprise tech landscape, bringing a diverse background in online content, ebooks, and language instruction.

Raconteur

Editor
Tom Dennis

Staff writers
Simon Chandler
Jay Lesada

Commercial production managers
Audrey Davey
Rachelle Kusi-Appouh
Bianca Schwartz

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

✉ @raconteur in raconteur-media 📺 @raconteur.stories

GLOBAL BUSINESS

Europe’s new market rules
force global business local

The rules of market access are changing in strategic sectors, forcing multinationals to weigh resilience and public support against the cost of duplicating capacity

Jay Lesada

The era of open, frictionless global market access is ending. From clean technology and energy to artificial intelligence and defence, European capitals are rewriting trade rules to enforce local production and sovereign resilience. For multinational executives, selling into the European market increasingly requires investing inside its borders.

This shift transforms corporate strategy. Boardrooms can no longer evaluate international expansion purely on cost efficiency or global scale. Chief executives and finance leaders must balance sovereign preferences, local content rules and regulatory barriers against the expense of duplicating operational capacity.

Nowhere is this clearer than in state procurement. In Poland, defence spending with domestic firms and Central European partners has quadrupled since 2022, reaching 30.4bn zlotys (£6.3bn) last year. Foreign suppliers aiming for government contracts must show a clear commitment to local investment. “If you want to sell in Poland, you have to invest in Poland,” deputy state assets minister Konrad Golota told Reuters.

Multinationals are responding. German defence group Rheinmetall is building production capacity and a maintenance hub in Poland. Czech manufacturer CSG has agreed to transfer propellant-production technology to domestic producer Mesko, expanding local manufacturing by acquiring component maker Domar MS. “Central Europe is undergoing a broader shift in how defence industrial capacity is built,” says Wojciech Grzonka, chief executive of CSG Polska.

The pattern extends across the continent. Germany’s rising defence spending draws overseas suppliers into local offices, factories and regional partnerships. In Romania, South Korea’s Hanwha Aerospace is constructing its first European production plant, aiming for up to 80% local industrial participation. Rheinmetall has also secured €5.7bn (£4.8bn) in Romanian state contracts, alongside plans for local investment, technology transfer and a supplier network involving more than 200 domestic businesses.

European Union legislation reinforces this operational pivot. Sébastien Maillard, associate fellow at Chatham House, describes the



growing push for European preference as a “change of paradigm compared to the traditional open-market approach” pursued by the bloc over recent decades. Member states, however, remain divided over how narrowly to define that preference.

The €150bn (£126bn) Security Action for Europe programme limits non-qualifying foreign component costs to 35% for eligible procurement. Sensitive defence systems face extra conditions, including whether contractors can modify equipment without restrictions imposed from abroad.

This policy direction extends far beyond defence. Under the Net-Zero Industry Act, public procurement and renewable-energy auctions must consider supply chain resilience and environmental sustainability alongside price. Contracting authorities can penalise bids dependent on single non-EU suppliers. Meanwhile, the proposed industrial accelerator act would introduce ‘Made in EU’ criteria and strict conditions on foreign investments in strategic sectors.

The UK takes a less prescriptive route. The Cabinet Office told central government bodies to take a more active role in shipbuilding, steel, artificial intelligence and

energy infrastructure, placing greater emphasis on security of supply and long-term resilience.

For global enterprises, building a European footprint improves market access and public support while securing supply chains against geopolitical disruption. The trade-off is high: duplicated manufacturing capacity, elevated operating costs, complex joint ventures and pressure to share technology or operational control.

Defence firms appear willing to pay that price. Foreign direct investment projects in European defence rose 84% in 2025, according to consultancy EY. It also found that 59% of defence executives plan to set up or expand European operations over the following year.

Yet this trend does not mark a wholesale retreat from global supply chains. A European Commission survey found that 70% of EU manufacturers with production outside the bloc have made no recent changes to those locations and plan none.

“European preference is meant for strategic sectors,” Maillard says, though disagreement remains over what counts as strategic. The trend is better described as selective localisation: placing specific capabilities

closer to the governments and customers that view them as essential to security.

Boardroom decision-makers face clear operational choices as Europe tightens its market rules. Chief executives and chief financial officers can no longer evaluate international expansion purely on cost efficiency. Capital allocation must account for regulatory compliance, local content thresholds and sovereign risk alongside operating margins.


Chief financial officers must re-evaluate the true cost of market entry. While duplicating manufacturing capacity in Europe increases capital expenditure, failing to meet localisation thresholds can wipe out addressable revenue in target markets. CFOs must balance the higher unit costs of regional production against the risk of losing access to state-backed procurement programmes.

Chief operating officers must manage the friction of regionalised supply chains. Setting up local plants often involves complex technology transfers, workforce training and shared management with domestic partners. Companies that successfully navigate these partnerships establish clear operational boundaries and protect core intellectual property while giving local authorities the domestic economic footprint they demand.

Chief technology officers and risk officers must evaluate data sovereignty and technology control. As European regulations tighten around sovereign cloud services, artificial intelligence infrastructure and software, multinationals must decouple regional data architectures from global parent networks to ensure compliance.

Finally, board members must define which parts of their product portfolio require local assembly and which can remain global. Selective localisation means identifying high-risk, heavily regulated sectors – such as defence, clean technology, steel and software infrastructure – where a local footprint is non-negotiable, while maintaining global supply chains for less sensitive commercial lines.

As European capitals harden their stance on industrial autonomy, the era of frictionless global procurement is fading. Multinationals that adapt quickly by treating local investment as a strategic asset, rather than an administrative burden, will secure a decisive advantage in Europe’s changing marketplace. ●



AI won’t transform the enterprise
unless leaders redesign it

AI delivers lasting value when organisations redesign operations, workforce capability, accountability and measurement around the business outcomes that matter

AI is now commonplace in corporate life. Yet while organisations have invested confidently in the technology, evidence it is improving performance remains harder to find.

New research from PA Consulting captures the gap. The average worker uses AI four times a day, yet one third of employees say its impact still lags the attention it receives. And more activity does not necessarily mean more value: 37% spend as long reviewing AI outputs as they would completing the work manually.

“The assumption is that using more AI always equals more value, therefore we only need to measure usage,” says Charlotte Townsend, people and change expert at PA. “But doing so doesn’t show you where it’s delivering impact.”

The organisations getting more from AI begin with a desired business outcome, not a deployment target. They work backwards from that result, redesigning the enterprise around it: the decisions people make, the roles and capabilities they need, where accountability sits and how capacity released by AI is reinvested.

Start with the destination

Each year, tens of millions of passengers pass through Amsterdam Airport Schiphol, one of the world’s busiest passenger hubs. A delay during aircraft turnaround can quickly spread beyond a single flight, disrupting gates, ground services and passenger connections.

Working with the airport, PA supported the development of Deep Turnaround, which uses AI technology to give operators earlier visibility of delays and emerging disruption. Instead of reacting shortly before departure, teams can intervene sooner, creating smoother passenger experiences. For the airport, it means better use of gates, stands and ground-service resources, while safeguarding operational resilience.

An earlier warning is useful only if someone can act on it. At Schiphol, the prediction changes when operators intervene and what they can prevent.

“You start from where you want to be,” says Rahul Harlalka, global digital lead – transport at PA. “At Schiphol, that meant understanding the impact of the delay before deciding where to intervene. You then work backwards and put the right interventions in place, so that each one contributes to the end goal.”

This means looking beyond the processes that are easiest to automate. Automating a simple task may improve a

dashboard without improving the business. Leaders must identify the decisions that constrain performance, then determine what needs to change around them. Doing so means redesigning roles, decision rights, incentives, governance and performance measures, not just introducing a faster tool.

Redesign capability with the work

Redesigning the flow of work also changes how people acquire the knowledge needed to exercise judgement within it. PA’s study of AI use across 2,000 private and public sector employees shows that senior leaders report substantially greater gains in productivity, capability and quality than those in junior, entry-level and non-managerial roles. Senior leaders are nearly twice as likely to say AI enables them to do things they previously could not, and around 1.7 times more likely to report quality improvements.

Experienced leaders can assess an AI output against knowledge built over years. Junior colleagues are still developing that frame of reference, and have historically done this through the routine research, drafting and analysis that AI is beginning to absorb.

If some of these routine tasks junior employees traditionally learn from are automated, organisations must deliberately create new pathways for employees to develop expertise through mentoring, observation, challenge and practice, rather than assuming development will happen on its own.

Senior leaders must also make their own standards more visible. “The best leaders make time to explain the difference between okay, good, and brilliant,” says Townsend. “If people have fewer opportunities to build judgement by doing the work themselves, you have to create other opportunities for them to develop it.”

Dependence may be the other side of familiarity. One fifth of employees feel overly dependent on AI, rising to 28% among frequent users, while heavy AI users are also more than 2.5 times as likely as light users to favour an AI answer over their own judgement.


This makes ongoing training as important as initial adoption. AI capabilities and limitations continue to evolve, and extensive use can change how employees perceive both the technology and their own judgement. Without regular refreshers, employees risk losing the ability to distinguish a genuinely useful output from one that merely sounds convincing.



“
Saving time is not the same as creating value. Organisations must make sure that any time saved is reinvested into high-value outcomes, including upskilling

AI’s value dividend is flowing disproportionately to senior leaders

Percentage of respondents agreeing with each statement



Make responsibility visible

AI produces an answer, not responsibility for how that answer is used. Although two-thirds of workers believe responsibility remains with the output owner, 34% feel less responsible for quality when AI is involved – and more than a quarter feel uncomfortable signing off work shaped by AI.

Leaders therefore need to define where human judgement is required, who owns the decision and what evidence is needed before work is approved. A faster decision is not necessarily a defensible one. Employees should be able to interrogate AI-assisted work and explain how the result was validated.

As Townsend puts it: “A team using AI selectively but making higher-quality decisions may outperform one with heavier usage and weaker judgement. The goal isn’t ubiquitous use. It’s capable and accountable use.”

Decide what saved time should buy

The same shift is needed in measurement. Pilots launched, licences activated and hours saved are easy to count but say little about whether work has improved.

“Saving time is not the same as creating value. Organisations must make sure that any time saved is reinvested into high-value outcomes, including upskilling,” says Warwick Goodall, global roads and mobility lead at PA.

Reinvestment should be designed into change from the outset. If leaders wait until time has been saved to decide what to do with it, efficiency may never become enterprise value. This principle has been applied in practice during large-scale Microsoft 365 Copilot

deployments. Working with PA, a leading European bank established a value measurement framework to support the expansion of users. Built around individual productivity and business results, the framework was designed to measure not only whether employees worked faster, but what value the organisation generated from the capacity that was freed up.

Adoption is only the first signal. Leaders also need evidence that behaviour has changed, outcomes have improved and the capacity released by AI has been put to better use. Otherwise, a successful rollout can coexist with stagnant performance.

Together, these signals show where progress is stalling. If adoption rises but outcomes do not, the obstacle may lie in the operating process, workforce capability, decision rights or the use of saved time.

Enterprise AI will not be won by the organisation with the most licences or the busiest users. The advantage will belong to those that redesign the enterprise around the outcomes they want: turning intelligence into better decisions, saved time into additional value, and widespread adoption into stronger performance.



Scan here to explore the Intelligent Enterprise



Bringing Ingenuity to Life.

Artificial Innocence

Existing statutes weren't built for autonomous systems, leaving executives to navigate the grey area between negligence, contract law, and common law precedent

Simon Chandler

It's a question most enterprises will have to answer sooner or later. Are they liable if they deploy an AI agent and it ends up doing something illegal? At least some may be hoping the answer to this is 'no,' judging by how frequently agents have gone 'rogue' in recent weeks. First there were exploits from OpenAI, Anthropic, Moonshot AI and Meta models, which had all basically absconded

from their test environments to hack external sites or platforms. And not long after, news broke that an Australian man had inadvertently hacked a gym's booking system, after he tasked a Claude-based agent with securing him a place in a pilates class. What these cases highlight is that AI agents can have a myopic focus on completing assigned tasks, sometimes acting in complete disregard

of explicit guardrails or proscriptions. Because they're probabilistic systems, predicting how they'll respond to prompts can never be an exact science. This leaves the door open to transgressions, and while Anthropic revealed that only three out of 141,006 of its evaluations resulted in "real-world incidents," even this tiny error rate could be too big if agents end up completing billions of actions per day.

In other words, organisations need to prepare themselves for the possibility that the agentic AIs they use will create problems. Not operational problems, but legal problems. And according to solicitors and legal experts, these problems have no easy answers. They will generally have to be decided on a case-by-case basis, and this is unlikely to change in the UK anytime soon, even if and when the Computer Misuse Act receives a planned update.

Aside from several cases in the United States involving Tesla vehicles, there are no recorded instances of anyone being prosecuted for the actions of an automated system. In California, a Tesla driver faced no contest in 2023 to two charges of vehicular manslaughter, while in 2024 Tesla itself settled a long-running lawsuit that blamed the EV manufacturer for the death of a Model X driver who fatally crashed into a barrier near San Francisco. In both cases, Autopilot had been activated, but the seemingly opposed outcomes underline the difficulties of deciding where ultimate culpability resides.

In the UK, there are no precedents as relevant as the above Tesla

“It is unclear whether liability law applies to both AI-embedded products and standalone AI software

accidents. And to complicate matters, solicitors explain that liability could vary considerably from case to case. Although one thing is certain: it won't be the AI agent that ends up with a conviction.

“AI has no legal personality under English law, so it cannot be held responsible as an independent entity,” says Brett Dixon, the Vice President of the Law Society of England and Wales. “Depending on the circumstances around deliberate harm and legal obligations, it could be the user, developer, or platform/service provider that could be held liable.”

Other experts share similar views. Rory Coutts, a cybersecurity lawyer at tech-focused legal firm Bird & Bird, suggests that AI agents going 'rogue' is a decidedly grey area. As such, the question of whether an offence has been committed “is not clear-cut,” and neither is the question of who exactly is the offender.

“Liability is fact-specific and would depend on the individual case,” he says. “The Computer Misuse Act criminalises unauthorised access to a computer, with liability turning on intent, knowledge and who caused the act — who directs the agent — and whether system-level controls were in place.” So any company or organisation that deliberately tasked an AI with accessing external systems and platforms would almost certainly be criminally liable.

A company that deployed an agent according to all available developer and provider guidelines, and that maintained reasonable guardrails, would surely not be held accountable. However, here's where things become murky, since it's not entirely sure how the UK's product liability laws would apply to LLM-based systems.

“The Law Society has called for clearer guidance on AI-related professional standards so further legislation or regulation may be needed.”


Dixon also notes that the UK Jurisdiction Taskforce's view is that existing laws, including negligence and contract law, “are generally capable of dealing with AI-related harms.” As noted in commentary recently published in The National Law Review by solicitors from Squire Patton Boggs, this means that existing English and Welsh common law is capable of deciding most cases in which an AI system errs.

This commentary also noted that one way for organisations to resolve the thorny question of liability is to draw up contracts with AI suppliers. As the authors wrote, “the message is clear: the best way to allocate responsibility, mitigate risk and manage liability between parties in connection with the acquisition or use of AI systems (as well as avoiding the time and expense of becoming involved in a precedent setting case) remains a comprehensive and well-drafted contract.”

With relevant contracts signed, English common law would have to step in only where disputes arise and gaps emerge. While time may tell whether there are material gaps in common law, there remains the possibility that the UK Government will eventually introduce specific legislation.

“The King's Speech in May already announced plans to update the Computer Misuse Act,” explains Coutts. “This would be relevant where 'ethical' or 'white hat' hacking is carried out to identify vulnerabilities and report them to system owners.”

However, Coutts acknowledges that such an update may cover only a small subset of cases where a 'white hat' hacker uses AI with benevolent intentions, and that it would not “clarify where legal responsibility lies in all cases.” Plus, there's no guarantee that the new Burnham government will adopt the plan, meaning it could be some time before the UK benefits from any new, targeted legislation that has a direct bearing on AI (mis)use. ●



IBM, 2026; SailPoint 2026; NeuralTrust, 2026

‘Companies must remain transparent about how they use AI and data’

As technological shifts reshape workplace culture, Sage chief people officer **Amanda Cusdin** explains how maintaining transparency and building early career pipelines ensures digital progress does not erode organisational values

Simon Chandler

Trust and transparency are watchwords for Amanda Cusdin. Having served as chief people officer at Sage for eight years, she recognises that honesty sits at the centre of building a thriving corporate culture. Creating an open environment where staff trust leadership and colleagues underpins her strategy.

“We are creating a culture where people trust each other, trust the colleagues around them and trust what we are doing and where we are going,” she explains. “That is really important to us.”

Trust brings practical internal and external benefits for Sage. The business reported an 11% year-on-year increase in total revenue for H1, driven by expanding AI and cloud divisions. Cusdin notes that Sage's priority on trust helps the company stand out in technology, aiding recruitment and customer acquisition. Prospective employees regularly tell recruiters that they want to join Sage because of its culture.

Trust alone does not account for Sage's growth or Cusdin's work leading its global HR function. The business has also benefited from a willingness to change early. This proactive approach shows across recruitment, training and internal operations, allowing the company to adapt before external pressures force change.

Cusdin's career extends beyond her tenure at Sage. She studied history at university, which gave her a broad perspective on workforce management. Her early roles at Novar and Invensys allowed her to manage HR across multiple departments and regional offices, building a clear understanding of employee needs across operational units.

Her work at Invensys included its rail signalling division, where she handled recruitment, training, development and organisational effectiveness for engineering teams. Operating in this environment taught her how HR initiatives directly support frontline technical teams.

Novar and Invensys were manufacturing and engineering companies producing physical goods. This environment encouraged a direct, practical HR management style. Vague HR concepts carried little weight with signalling engineers who preferred straightforward answers. Direct roles dealing with managers and employees helped her assess how HR initiatives made a practical difference to daily operations.

Schneider Electric acquired Invensys in 2013. Soon after, Sage appointed Cusdin as vice-president of human resources. She accepted the role, joining at the start of Sage's move to cloud computing. The appointment gave her an opportunity to support the business as it transitioned into a single operating model.

Sage was moving from a federated business structure into a single operating model. Cusdin provided HR support to several key functions undergoing transition, including marketing, finance and product development. Spending her early years at the business building a practical understanding of its operations clarified what her people strategy needed to address when she became chief people officer.

Since her appointment as chief people officer, Sage has completed both a cloud transformation and a Software-as-a-Service transition. This experience positions Cusdin and Sage to manage an ongoing AI



“HR teams must focus on the business problems they solve, not just the HR processes they deliver

transformation. Previous internal evolution has made the workforce adaptable to external disruption. Even without external market shifts, Sage would have evolved its candidate profile, skill requirements and internal culture.

Sage's recruitment profile changed during its move to a unified model, requiring employees capable of working across multiple markets while developing local products. Recruitment criteria are evolving again to meet the demands of AI, which calls for distinct technical and analytical skills. Sage focuses on recruiting and upskilling talent to apply AI effectively while preserving customer and employee trust.

Artificial intelligence has introduced uncertainty into business operations, as organisations adopt new tools without fully understanding long-term consequences. In this

climate, Sage places greater emphasis on trust and transparency, addressing enterprise concerns about control and autonomy. Companies worry whether adopting AI tools might lead to a loss of operational oversight.

“Trust and technology are two words that go together, but we really, truly believe that we want to be a technology company that people trust,” Cusdin says. “That means transparency about how we use AI and how we use data.”

Sage ensures customers retain ultimate decision-making control over AI capabilities embedded within software products. Customers choose whether to use AI features. This transparent approach applies internally too. Recruiters maintain clear communication with prospective candidates, who regularly confirm that workplace culture matches initial recruitment descriptions. Candidate expectations align with actual employee experiences.

Sage employs 11,500 people globally, up from 11,170 at the end of 2025 and 10,789 the year before. The business uses external hiring alongside internal upskilling to manage AI workloads, equipping staff with tools to increase capacity and overall output. Cusdin aims to give employees the resources

needed to use AI effectively in their daily responsibilities.


Maintaining a balance of early-career joiners and experienced staff remains a key operational goal. Graduates and early-career employees bring fresh perspectives and maintain long-term leadership pipelines. Rather than scaling back entry-level routes, Sage is expanding its graduate and apprentice schemes. The company will welcome 500 early-career joiners later this year, aiming to make early careers 10% of its global workforce by 2028.

While roles will adapt as technology advances, people will remain central to Sage's operations. Cusdin considers empathy and active listening essential qualities for HR leadership. Chief people officers must aim to create environments where employees succeed, remaining humble enough to challenge existing practices as the commercial environment changes.

HR functions must focus on business outcomes rather than rigid processes. Working across departmental boundaries enables HR teams to solve core business problems rather than delivering isolated HR initiatives. By aligning people management closely with strategic goals, HR leaders ensure corporate culture supports long-term commercial growth. ●


WHERE THE C-SUITE IS HARVESTING REAL GAINS

Percentage of enterprise leaders identifying key structural, talent, and governance barriers




WHERE THE S-SUITE IS FOCUSING IN Q4

Based on executive responses measuring tangible business outcomes from strategic deployments




BOARDROOM FRICTION FACTORS

Percentage of enterprise leaders identifying key structural, talent, and governance barriers



STRATEGIC INTEGRATION & REALITY CHECKS

How executive boardrooms are shifting from speculative pilots to disciplined scaling



THE 2026 C-SUITE AGENDA

Enterprise leaders are pivoting from chaotic AI experimentation to hard-nosed execution, yet structural inertia remains a primary bottleneck. While boardroom optimism around tech-driven productivity rises, mounting geopolitical risk, severe skill deficits, and legacy operating models are forcing executive teams to drastically re-engineer their growth strategies.



Why adding more AI tools will not make work flow

Businesses have been quick to adopt a host of AI tools, but many have struggled to turn them into meaningful outcomes. To realise the technology’s potential, they need intelligent workflows that connect AI to how work gets done across the organisation


For many businesses, AI adoption has happened one tool at a time. Employees have gained new ways to automate or accelerate individual tasks, but those tools often sit within the same processes, systems, and organisational boundaries that were already in place. The challenge now is to bring more structure to how AI is deployed across the business. Data from Gartner, based on a survey of 1,300 leaders, found that just 22% of enterprises have successfully scaled AI across multiple business departments. Meanwhile, just 37% of 1,700 leaders surveyed by McKinsey said AI had increased profitability. Fragmentation also creates governance and security challenges. Leaders need human oversight and clear controls over AI access to sensitive business and customer data. Leaders must move beyond AI tools that shave time off individual tasks and use the technology to improve workflows across the business. That creates opportunities for greater efficiency and profitability. That shift matters as enterprise AI evolves from language models that help employees complete individual tasks to agents that can act across

multiple systems and stages of work. Their value depends on how effectively those actions connect into an end-to-end workflow built around a clear business outcome. Why businesses need intelligent workflows OneAdvanced, a provider of AI-powered, sector-focused SaaS software, is helping organisations navigate that challenge. Marko Perisic, the firm’s Chief Product Officer, says the current use of AI fails to create the intelligent, enterprise-wide change needed to drive tangible outcomes. “Most organisations, unfortunately, go down the path of getting an AI subscription of some sort, giving it to employees, and assuming that it will improve how they work and how they produce output,” he says. “But when you just apply AI on top of archaic workflows and sub-optimal ways of working, all you get is an AI supercharged mess. In a lot of cases, this creates more problems and compliance issues than positive outcomes.” Organisations seeing the greatest benefits from AI are using it to rethink existing workflows or design new ones that connect work across the business. Perisic says doing this requires a

deliberate approach. “You need to go down a path of conscious, purposeful, and structured transformation,” he says. “You need to rethink those workflows from the ground up and ask: ‘What do we actually want to try to accomplish?’ In the world of AI, perhaps you don’t need Step One, Three, and Seven of that workflow. Maybe you only need Two and Four, but they need to be very different from what they used to be. When you do that, that’s when you get the real benefits of AI.” Designing persona-first workflows OneAdvanced works with organisations to create persona-first workflows. This means designing the way work gets done around the specific needs, responsibilities, and context of the person using the software. The aim is to make technology fit the way people work, instead of forcing employees to adapt their work to a proliferation of AI tools. The process begins by identifying a specific role, such as a Financial Controller or Chief Marketing Officer, and mapping the key workflows involved. OneAdvanced then assesses the potential return in time, efficiency, or cost, and prioritises those with the greatest value for the customer.

“When you apply AI on top of archaic workflows and sub-optimal ways of working, all you get is an AI supercharged mess

Businesses are still targeting AI for productivity first



“It’s not enough to just analyse one step of a workflow,” says Perisic. “You need to really understand what can be done better using AI, what can be done faster using AI, how to do it in a safe and compliant way, how to do it in a way where it actually fits with the rest of the organisational workflows that maybe aren’t AI-powered at the same time, and where the interface points.”

An intelligent workflow in action

OneAdvanced IQ is designed to put this approach into practice, connecting people, data, workflows, and AI in one system of work. A law firm recently partnered with OneAdvanced to use IQ to automate client onboarding. The process was time-consuming and non-billable, but essential for compliance, including anti-money laundering checks. Automating parts of the workflow allowed legal experts to focus on higher-value work. In practice, that means connecting the different stages of onboarding rather than automating a single check in isolation. Once a new-client request triggers the workflow, AI agents can gather the required information, carry out the relevant checks, flag exceptions, and return the evidence for review. Where a result is ambiguous or requires professional judgement, the agent will escalate it to the appropriate person rather than making the decision autonomously. “It’s had a huge transformational impact for these individuals,” says Perisic. “They can tell the AI about the client they want to onboard and the checks they need to make, and then AI agents will run that workflow in the background. At the end, it will report back with a list of completed checks, and then the client can be successfully onboarded. Legal professionals can then do the work they’re trained for, minimising administrative tasks.” But while AI can be used to carry out workflows autonomously, human

oversight is still critical to ensure the security and governance of sensitive customer data. Perisic says organisations should begin by limiting the scope of AI. “You should start super narrow with very limited permission sets,” he explains. “Then you can consciously expand that with controls, so that people understand that agents have access to certain data sets or have permission to send emails or carry out other tasks.” Permissions form one part of a wider control framework. Organisations need visibility into where sensitive data is held and processed. They should also be able to understand and explain how AI reaches decisions or recommendations, with clear checkpoints for human review, approval, or intervention. Safeguards should reflect the potential impact on customers, patients, or citizens. Organisations can start with a tightly defined, high-value workflow, establish how it performs today, and automate specific stages while retaining human controls. They can then measure the impact, identify where exceptions arise, and gradually expand the agent’s remit. This limits disruption and gives employees time to adapt. The next phase of enterprise AI will depend less on how many tools organisations deploy and more on how effectively they connect AI to the work that people already do.



For more information please visit oneadvanced.com

oneAdvanced

Q&A

Five questions on AI agents

Marko Perisic, Chief Product Officer at OneAdvanced, explains how leaders can introduce AI agents with clear goals, controls and human oversight



Q Before deciding which AI model or fleet of agents to use, what business outcomes should leaders define?

A Ultimately, every business wants to increase profitability, but that’s a lagging indicator. You need to work backwards from there to identify the leading indicators of success. For example, if you want to increase profitability, do you want to do that by reducing costs? Or do you want to do that by increasing growth? That will be one decision you need to make. If it’s increasing growth, how do you increase growth? You need to increase your win rate. How do you increase your win rate? You do so by competing, offering better products and services in that particular target market segment. And you work yourself backwards. Eventually, you’ll arrive at the leading indicators of success. And that could be increased usage, better NPS from customers, better customer sentiment that eventually is going to lead down the chain to positive economic indicators of performance of the business.

Q How can businesses measure whether AI is actually improving a workflow versus just improving the ability to perform individual tasks?

A It comes back to lagging and leading indicators. If AI is not changing the ultimate measures of a successful business (profitability) either by reducing costs or increasing growth, then it’s not really useful. Now, it may reduce the time that employees need to spend doing work or learning how to do work. So that will increase employee satisfaction, which in turn can also benefit the economic performance of the business. You get better employee retention. You get better engagement, and better-engaged employees produce better work. Ultimately, it really comes down to growth and profitability. So, I think you just need to keep that in mind when you look at it. Because otherwise, as I said, you’re just burning tokens and everybody gets busy, but you’re not producing the output you want to produce. **Q What key boundaries should businesses set if they’re using agents for the first time?**


A First of all, I’d make sure you have a very purposeful and thoughtfully designed security posture in the organisation and a dedicated team that’s managing that. You need to understand where your critical assets are, where your critical data is, who has permission to access it, and how and when they access it. You need to narrow those permissions down as much as possible. One approach is to adopt zero-based security. This is basically no permissions, with permissions only granted by exception and acceptance. You need to start narrow and expand rather than give everyone access and mitigate when it goes wrong. Because when it goes wrong, it’s already too late. And it’s harder to untangle. **Q Where should human judgment and oversight remain?**

A I think it needs to be measured by AI’s impact on other human lives. All businesses have an impact on human lives. If my pizza delivery is half an hour late because some agent screwed up, it’s not the worst thing in the world. But if my medical prescription is wrong

“Giving employees AI, hoping for the best, and mitigating problems afterwards isn’t the way to go. You need to start with a structured view

because an agent prescribed the wrong medicine based on a diagnosis it produced, then that could be fatal. And that needs to be very critically controlled and addressed. The severity of the potential impact on people’s lives should determine the level of human control. And also, for that very reason, there are laws and compliance rules around some of these things. What can be used in certain industries and what cannot. We adhere to those very strictly and thoroughly. I think there’s going to be more and more focus on that in the industry because there’s a lot of AI sprawl. If it’s not controlled, it could be damaging. **Q In the first 90 days, what should I measure to understand whether AI implementation has been a success?**

A Be clear about your lagging indicators and the leading indicators you’re going to measure. Consider the severity of the potential impact on people, how many controls the workflow needs, and who owns them. Document this in a structured, centralised place that the right people can access. That gives you a foundation for deciding how to move forward. Giving employees AI, hoping for the best, and mitigating problems afterwards isn’t the way to go. You need to start with a structured view.



INDUSTRY INSIGHTS

Beyond the hype: where UK deep tech is building sovereign capability

With capital shifting from software hype to physical infrastructure, three UK deep-tech leaders explain why hard tech and sovereign supply chains hold the key to long-term value

Simon Chandler & Tom Dennis

Building physical technology at scale presents a starkly different challenge from software development. While digital products benefit from rapid iteration and low marginal costs, hard-tech ventures must bridge the gap between

complex physics and heavy infrastructure—a transition often called the “valley of death”. Succeeding in this environment requires long-term capital, manufacturing discipline, and a realistic approach to global supply chains. As market dynamics shift

focus toward operational resilience and sovereign capability, leading UK scale-ups demonstrate that the true test of deep-tech leadership lies not in laboratory breakthroughs, but in industrialising complex physical assets for global markets.



Quantum Hardware & Advanced Semiconductor Manufacturing

James Palles-Dimmock

CEO, Quantum Motion

While the rapid expansion of artificial intelligence dominates headlines, it has also reminded institutional investors that entirely new compute paradigms create immense value. Classical computing accelerated because specialized hardware like GPUs made complex calculations dramatically cheaper. Quantum computing addresses a parallel premise: what happens when an entirely different class of computation becomes accessible at scale? However, because quantum hardware represents a longer-term proposition, capital markets have grown increasingly discriminating. Investors are filtering out laboratory demonstrations, focusing instead on whether a deep-tech architecture possesses a verifiable, cost-effective path toward manufacturing the millions of physical qubits required for meaningful real-world commercial applications.

To address this scaling requirement, Quantum Motion deliberately avoided inventing a novel manufacturing ecosystem. The fundamental engineering question was not merely how to create a high-performing qubit in a laboratory, but what existing industrial base could reliably produce millions of near-identical components. The answer lies in standard silicon transistor and CMOS semiconductor manufacturing – the infrastructure powering global smartphone and computing chips. By adapting established CMOS processes, the company leverages decades of prior industrial investment, high device density, and global foundry capacity rather than attempting to recreate a semiconductor industry from scratch, allowing the platform to scale rapidly.

This thesis has transitioned from physics to system-level execution.

Key operational milestones include deploying a full-stack silicon CMOS quantum computer at the UK’s National Quantum Computing Centre (NQCC), advancing into Stage B of DARPA’s Quantum Benchmarking Initiative, and expanding foundry relationships with GlobalFoundries and international partners. This trajectory underpinned a \$160-million Series C funding round, supported by specialist tech investors like DCVC and Kembara alongside the British Business Bank. This investor syndicate reinforces a central thesis: the quantum race will ultimately be won through manufacturing discipline and industrial scale rather than laboratory promises.

From a leadership perspective, scaling hard tech requires a corporate culture rooted in scientific falsification. Executives must maintain a clear destination while encouraging cross-disciplinary teams – spanning physicists, chip designers, and software engineers – to test and disprove underlying assumptions quickly. Treating new empirical evidence as a reason to change direction is a core operational strength. Furthermore, while geopolitics increasingly frames semiconductors and quantum as sovereign technologies, true resilience requires international supply chain integration rather than isolation. The ultimate objective is not building a few billion-dollar machines for nation-states, but making quantum compute abundant and accessible for global commercial application.



The quantum race will ultimately be won through manufacturing discipline rather than laboratory demonstrations

Industrial Automation & Supply Chain Robotics

Rebecca Marsden

SVP commercial and strategic partnerships, Oxa

Navigating macroeconomic headwinds has made raising growth capital increasingly challenging for deep-tech businesses with extended investment horizons. While artificial intelligence has attracted significant market liquidity, the primary challenge for scale-ups is differentiation. Investors are evaluating automated systems with heightened rigor, demanding defensible technology, a credible route to market, and a clear transition from technical capability to measurable commercial value. Founded in 2014, Oxa differentiates itself from recent market entrants by drawing on over a decade of operational experience developing and deploying mature autonomous software in safety-critical industrial environments where software reliability directly impacts operational continuity.

Rather than pursuing unconstrained passenger transport applications, Oxa focuses its strategy specifically on Industrial Mobile Autonomy. Automating repetitive vehicle movements within structured industrial environments – such as ports, airports, logistics hubs and manufacturing facilities – targets defined sectors where operational budgets already exist. In these settings, autonomous vehicle software delivers immediate, measurable gains in safety, fleet productivity and supply chain resilience. The economic case for industrial customers is clear and near-term, offering a far more direct route to commercial return than unconstrained autonomous vehicles operating on public roads, delivering sustainable value for operators.

This focused commercial strategy is validated by Oxa’s recent \$103-million Series D first close. The investment syndicate combined strategic partners such as NVentures (NVIDIA’s venture arm) and bp Ventures with institutional backing from IP Group, Hostplus, and a \$50-million commitment from the UK National Wealth Fund. This blend of strategic and public capital reflects strong investor confidence in physical AI to translate software leadership into industrial productivity. Further-

more, commercial rollout is advancing through strategic joint ventures, such as SHIFFT with the Dubai Future Foundation, focused on deploying autonomous logistics across port and aviation environments before the end of 2027.

In a crowded investment environment, executive leadership requires proving that physical AI addresses identifiable operational challenges with a clear path to scale. Beyond technical sophistication, investors and industrial clients require evidence of market demand, established safety track records, and sustainable unit economics. By deploying mature autonomy into structured industrial workflows, scale-ups can capture existing corporate budgets while mitigating deployment risk. Ultimately, long-term value in robotics and physical automation is created when technology transitions out of experimental pilots and into daily commercial operation, establishing defensible, highly resilient logistics infrastructure for global markets.



Clean Tech & Sovereign Critical Minerals Refining

Dr. Christian Marston

President, COO & co-founder, Altium

Achieving net-zero targets requires widespread electrification across transport and energy grids, driving massive demand for critical battery minerals like lithium, nickel, cobalt and graphite. However, global supply chains are heavily concentrated, with China controlling the vast majority of refining capacity – including 96% of graphite processing. As geopolitical tensions and tariffs prompt nations to protect industrial resilience, relying on overseas refining creates significant economic exposure. In the UK, with over two million electric vehicles on the road, end-of-life batteries represent a strategic asset – an “urban ore body” containing high concentrations of critical minerals ready for domestic recovery.

To capture this resource, Altium developed eco-friendly processing technologies capable of recovering over 95% of critical minerals from battery waste indefinitely. This circular approach produces materials with a 74% lower carbon footprint than traditional refining, reducing a new EV’s total carbon footprint by 20%. To manage infrastructure execution risk, Altium follows a controlled, staged scaling model. Progressing methodically from bench testing (ACT 1) to pilot operations (ACT 2) in Devon, the company is commissioning its first commercial intermediate facility (ACT 3) in Plymouth to process waste from 24,000 EVs annually, laying the foundation for a mega-scale refinery handling 150,000 EVs per year.

This disciplined approach avoids a common pitfall in clean tech: hyper-scaling too quickly on early capital before fully proving operational processes. By managing scale-up risk in stages, Altium secures patient capital from global industrial part-

ners, including Series A funding from SQM – the world’s largest lithium producer – and Series B investment from Marubeni and Mizuho Bank. This capital is backed by an £18.5-million grant through the UK government’s Advanced Propulsion Centre (APC) Drive 35 programme, covering 50% of the capital and operational costs for the ACT 3 plant, providing strong public validation for sovereign refining infrastructure and domestic growth.

With EU regulations requiring minimum recycled content in EV batteries by 2030 and trade barriers increasing, domestic refining creates strong competitive moats. As the only European producer capable of supplying 100% sovereign, non-Chinese cathode and anode active materials, Altium provides critical supply chain security for automotive and defense applications where overseas exposure is unacceptable. Combining proprietary clean technology with deep commodity expertise and strategic grant co-funding enables hard-tech scale-ups to navigate capital-intensive infrastructure development, bridge the “valley of death,” and establish resilient industrial assets. ●



Over-capitalisation can force premature expansion. De-risking physical scale-up requires a disciplined progression from pilot facilities to commercial refineries

Q&A

The hidden barrier to AI success

New research from Expereo and IDC suggests that while organisations are racing to adopt AI, few have built the network foundations needed to realise its full value. **Ben Elms**, CEO of Expereo, explains why network readiness is becoming the difference between AI ambition and AI success



Q AI investment is booming, but your Enterprise Horizons 2026 research suggests many organisations aren’t seeing the returns they expected. What’s going wrong?

A The conversation around AI has become dominated by what’s possible, rather than what’s practical.

Most organisations understand the opportunity. They’re investing in new models, applications and use cases because they recognise AI’s potential to transform their business. The problem is that many are scaling ambition faster than capability, without fully understanding whether their networks can support the demands AI places on data and users.

Many organisations have yet to build the foundations needed to deliver consistent returns from AI investment, with challenges around data quality, skills, governance, infrastructure and visibility often holding back ROI rather than the technology itself. As AI workloads become more distributed, organisations need a clear understanding of network performance and potential bottlenecks.

AI success starts with resilient connectivity, end-to-end network visibility and the ability to deliver consistent performance at scale.

Q Why do many organisations underestimate the role that network infrastructure plays in AI success?

A Networks have traditionally been viewed as operational infrastructure rather than a strategic business capability. That mindset doesn’t work any more.

AI is fundamentally changing how organisations think about data. Workloads are increasingly distributed, real-time and data-intensive, so network performance becomes a direct enabler of business outcomes.

Our research found that inadequate networks are among the leading reasons AI failed to deliver expected

returns. Conversely, the organisations exceeding expectations consistently cited strong network performance as one of the biggest contributors to success. That tells us that infrastructure now determines whether AI investments succeed or not.

Q How are AI workloads changing the demands placed on enterprise networks?

A Organisations are increasingly moving large volumes of data between cloud environments, edge locations and distributed teams while expecting AI services to respond almost instantly. That places far greater emphasis on bandwidth, resilience and flexibility, which many existing networks aren’t designed to deliver.

The companies using AI in the most advanced ways already have significantly more capable networks. They understand that preparing for AI success depends on building connectivity structures that are agile, resilient and able to adapt as business needs evolve.

Q What separates organisations that are successfully scaling AI from those that are struggling?

A The biggest difference I see lies in executional discipline. Successful organisations don’t treat AI as a collection of isolated experiments. They invest across the foundations – strengthening data quality, building internal skills, establishing governance and ensuring networks are capable of supporting AI at scale.

Our research reinforces this. The firms where AI ROI exceeded expectations were significantly more likely to have strong in-house expertise, high-quality data and robust network performance. Whereas those struggling are far more likely to cite poor data quality, skills shortages and infrastructure limitations. The winners aren’t necessarily investing more in AI. They’re preparing better for it.

Q Why is end-to-end network visibility becoming such a critical capability in the age of AI?

A As AI workloads span multiple cloud environments, edge locations and third-party providers, organisations need far greater visibility into how applications and data are performing across their networks. Without that insight, it’s next to impossible to identify bottlenecks, resolve performance issues or understand where resilience can improve.

Better visibility allows organisations to shift from reacting to problems to proactively strengthening performance before potential issues impact AI initiatives. That’s essential when AI applications might be supporting customer experiences or business-critical decisions where downtime has immediate commercial consequences.

Q If business leaders are planning their next wave of AI investment, what should they prioritise to maximise long-term value?

A You need to think beyond AI itself. The businesses creating sustainable value won’t necessarily be those deploying the largest models or launching the greatest number of pilots. They’ll be the ones investing in the capabilities that allow AI to scale safely and effectively over time.

My biggest piece of advice would be to ensure you have the right team, skills and governance structures in place – backed up by a resilient, secure and flexible network infrastructure.



Scan the code to read the Enterprise Horizons 2026 report

expereo
faster to the future



oneAdvanced

Make work flow

Turn complexity into clarity

One trusted, intelligent system
of work, connecting people,
processes, data and AI.



**Scan to
explore IQ**