softcat

Business Tech Report 26/27

Contents

03 Welcome

04 Introduction

05 Two perspectives, one technology agenda

06 Cost control and budgeting in a world of uncertainty

07 "That system's a bit special" IT standardisation and rationalisation

08 Artificial Intelligence is no longer a future consideration

09 From AI experimentation to execution

10 AI isn't about the GPU or the model. It's about the data

11 Good vibes only

12 AI's workspace impact could be right under our noses

13 How do we build trust in public sector AI?

14 Agentic AI: Financial services' next operating model

15 Innovation has to earn its keep

16 Tokenomics and the cost of AI across cloud and platforms

17 The mic drop moment

18 Understand before you accelerate

19 Talk data to me

20 Following the breadcrumb trail to answer: Where is the work?

21 Networking at machine speed

22 Trust must be designed in

23 How data and AI sovereignty will shape your strategy

24 Can you trust AI?

25 How do you take control of AI agent permissions?

26 Secure networking needs to be built in, not bolted on

27 Post-quantum cryptography: Why 2027 is the year to get started

28 Keeping critical public services running

29 Adoption is the differentiator

30 AI agents and human expertise: Better together

31 How should work happen today?

32 The future of commercial operations

33 Global standards meet local realities: Rethinking multinational tech procurement

34 How MAT is reshaping public sector procurement

35 How AI is changing bids

36 Summary

Welcome to the Softcat Business Tech report 26/27

"Core IT services are more important than ever before. Getting them right opens the door to successful AI deployments. This year's Business Tech Report gives us a clear insight into what your priorities are right now: you told us that agentic, workspace, controlling costs, and cyber security are all at the top of your list, and we're listening.

Softcat will continue to be right beside you, supporting your organisation and making these priorities more productive, more manageable and more successful."



Graham Charlton
CEO, Softcat

AI-first, foundation-led: Why core IT still matters most

Core IT services, like cloud, devices, networks, cyber security, and data remain fundamental to progress.

In 2027, the challenge is deciding where technology can make the biggest difference, then making those changes at a pace the business can sustain.

What our survey says

This year's Softcat Business Tech Report draws on insights from **6,671 customer responses across 4,279 organisations** captured through our annual Customer Experience Survey. Looking beyond AI, the report explores the core technology areas that remain critical to so many organisations, while examining how cost, security and risk may influence the pace of innovation.

'AI-first', but not 'AI-only'

AI – in particular, AI agents – has now moved to the top of the agenda. Organisations are now tasked with moving from experimentation to an operating model where AI agents deliver lasting value.

To realise value from AI, organisations need resilient fundamentals: infrastructure, adaptable security, clear governance, trusted data and connected user experiences. In the era of AI, core IT is more critical than ever.

> "To realise value from AI, organisations need resilient fundamentals."

Softcat can help you set those priorities, strengthen the technology you already rely on and decide where new capabilities, including AI, can deliver value.

I hope the insights in this report help you make those decisions with confidence.



Dean Gardner
Technology Director

Top 5 technical priorities from the 2026 Customer Experience Survey

- **41%** AI agents
- **37%** End-user devices and software
- **32%** AI orchestration and integration
- **32%** Cost control and budgeting
- **31%** Cloud security

Two perspectives, one technology agenda

Every year we ask our customers what matters most, and every year the leadership and technical answers come back looking like two different surveys.

This year **866** leaders and **4,013** technical contacts replied, and the gap between them is the most useful thing in the data.

Leaders put maximising the impact of the technology they already own at the top, with using technology for innovation a whisker behind. Technical teams put cost control and budgeting first, followed by IT standardisation and rationalisation. Governance, risk and compliance sits third on both lists, which may be the only thing the two groups have agreed on since the hot-desking policy.

Read them side by side and they stop being different answers. “Get more from what we have” and “work out what we are spending, and on what” are the same instruction written in two dialects. One is the board asking for the outcome. The other is the engine room describing what it will take to get there. Efficiency, business continuity and technology sourcing fill out the rest of both top fives, and not one of them is achievable while nobody can say with confidence what the estate costs or how many of everything it contains.

James Seaman
Head of Advisory Services

Top 5 strategic priorities from the 2026 Customer Experience Survey

Leadership contacts		Technical contacts	
63%	Maximising the impact of your current technology	52%	Cost control and budgeting
61%	Leveraging technology to drive innovation	42%	IT standardisation and rationalisation
53%	Governance, risk and compliance	37%	Governance, risk and compliance
47%	Efficiency challenges	35%	Business continuity
42%	Using technology to minimise business risk	29%	Technical sourcing and procurement

Cost control and budgeting in a world of uncertainty

Understand what the money buys, then make deliberate decisions about where it goes.

Most technology budgets are built the same way: take last year's number, add the uplifts the vendors have warned you about, subtract whatever finance has asked for, and argue about the gap in a meeting that overruns.

Nobody in that meeting is looking at the estate. They are looking at a spreadsheet of the estate, which is a different thing.

Our 2026 Customer Experience Survey tells us that Softcat customers see cost control as a priority for leadership teams. I don't read that as a wish to spend less. I read it as a wish to know what the money is buying, so that choosing becomes possible.

Treat run costs as a choice

A large share of many technology budgets is spent on running existing services, and that spend is often treated as fixed. It isn't. It is a stack of contracts, each signed by someone under pressure, each with a renewal date, an indexation clause and a usage profile that nobody has checked since go-live. A renewal is an important opportunity to negotiate with a supplier. Let it roll over without review and you may miss an opportunity to negotiate. That is a spending decision too, even if you did not write it down.

> **"When the board asks for £2m for AI, the CIO can point to where it comes from instead of asking for more."**

Build the budget from services upwards

Three habits separate the organisations that control cost from the ones that discuss it. A baseline that shows what each service costs in full, including the people, the support and the risk, rather than the invoice alone. A renewal calendar treated as a decision calendar, with the work starting a year out rather than six weeks. And a budget built from the services upwards, so when the board asks for £2m for AI, the CIO can point to where it comes from instead of asking for more.

Prepare for the next year

Expect 2027 to test this. Vendor price rises and AI consumption could put further pressure on run budgets, particularly where that usage has not been planned for.



James Seaman

Head of Advisory Services

“That system’s a bit special” IT standardisation and rationalisation

Cutting out the complexity creates a stronger foundation for automation and AI.

Walk any estate with the people who run it and you will hear the same phrase within the hour: “that one’s a bit special”. The finance system customised in 2014 by a contractor nobody can trace. The second collaboration platform, kept because one department prefers it. Four backup products. Three ways to log in. Each was a reasonable decision at the time. Together, they can make change slower and harder to manage.

Complexity never appears as a line in the budget. It appears as the 14-month project that should have taken six. The security incident that spreads because nothing is configured the same way twice. The AI pilot that stalls because the data lives in nine places and agrees with itself in none of them. The recruitment problem, because you need people who understand all of it, and those people are retiring.

Make simplification part of innovation

Technical teams put standardisation and rationalisation near the top of their priorities in our recent Customer Experience Survey. Leadership teams put innovation instead. I would gently suggest these are the same priorities wearing different hats. Standardisation makes automation easier, and a clear view of the estate gives AI programmes a stronger starting point.

> “Rationalisation is not a cull. Culls are quick and usually wrong.”

Decide what stays and what goes

Rationalisation is not a cull. Culls are quick and usually wrong. It is a scored decision about every application, platform and supplier: what outcome it serves, what condition it is in, what it costs to keep and what it would take to retire. Keep, consolidate, replace, retire. Then standards, which should be a short set of defaults rather than a two-hundred-page policy, with an exceptions process that carries a visible price.

“A bit special” becomes a conscious cost instead of an accident. Measure progress in counts: applications, suppliers, platforms, identities. Lower numbers, fewer surprises.

Do the groundwork deliberately

My prediction for 2027 is that this stops being housekeeping. Many boards want AI, and AI programmes may find themselves doing rationalisation under another name. Better to do it on purpose, and earlier.



James Seaman

Head of Advisory Services

Artificial intelligence is no longer a future consideration

It is here, already reshaping industries, redefining competitive advantage, and driving change across every aspect of the enterprise.

From business models and customer experiences to operational processes, investment priorities, workforce productivity, and risk management, AI is challenging long-held assumptions about how organisations create value and deliver results.

This presents both significant opportunity and considerable complexity. While the potential benefits are well understood, the path forward is rarely straightforward. Every organisation starts from a different position, faces a unique set of constraints, and has its own ambitions for transformation.

The question has evolved from “Where do we start?” to “Where do we go next?” and, just as importantly, “How do we move forward safely, securely, and at pace?”

Success demands more than access to technology alone. It requires the right balance of strategy, data, infrastructure, security, workplace transformation, cloud, and AI capabilities working together to create the conditions for sustainable innovation. Organisations must navigate uncertainty around governance, cost, adoption, and emerging technologies, while continuing to deliver measurable business outcomes.

Data remains the foundation of this journey, while AI tooling serves as the catalyst that transforms it into meaningful insight and tangible impact. Yet leaders are increasingly challenged with unlocking its full value while maintaining control, trust, and resilience in an ever-changing landscape.

This report explores the opportunities, challenges, and practical considerations facing organisations today. More importantly, it highlights the importance of partnership, bringing together the expertise, experience, and capabilities needed to turn AI ambition into meaningful business value and long-term competitive advantage.

> **“The key question is now ‘where can AI agents create real business impact?’”**



Andy Crossley
Chief Technology Officer, Oakland

From AI experimentation to execution

Technology leaders are being asked to drive transformation even faster, all while exercising greater control over how they do it.

Innovation, efficiency, resilience and governance are no longer separate conversations. Increasingly, they are shaping the same decisions.

Nowhere is that clearer than AI. The conversation is moving beyond what AI can do to what happens when it becomes part of everyday work. The opportunity is significant, but so are the questions around value, cost, trust and execution. That is where this year's story begins.

Adam Harding

Head of the Office of the Chief Technology Officers (OCTO)

AI isn't about the GPU or the model. It's about the data.

Moving from AI experimentation into production starts with the right data foundations.

Most organisations already have one or more AI workloads deployed and in production. But many projects are either hitting major roadblocks or failing to deliver the expected ROI. Solving these problems will be a defining challenge for 2027. The work starts not with models or GPUs, but with data.

In our conversations with organisations, the barrier to production is rarely the AI itself. Data quality, governance and ownership often have not kept pace with ambition. Without a strong data and analytics foundation, even a well-chosen model and a skilled team can struggle to deliver value.

Build the foundations before choosing the platform

Finding the right AI blueprint remains a major challenge for many customers. A successful pilot can show what is possible, but turning it into a production service calls for a clear understanding of the data, the workload and the platform on which to run it.

The organisations moving fastest are not necessarily using the most advanced models. They are the ones that addressed their data foundations first, then chose infrastructure and tools with the flexibility to adapt as requirements evolve. That gives them room to change without tying every decision to a single platform.

> "The organisations moving fastest are not necessarily using the most advanced models."

Match infrastructure to the workload

Softcat's prediction for 2027 is that infrastructure decisions will increasingly be driven by the needs of production AI workloads. Vendor preference will carry less weight than how well a platform fits the task.

NVIDIA remains central to AI compute, but how that compute is deployed matters as much as the GPU itself. Public cloud providers offer speed and elasticity for experimentation, while private on-premises options from Dell, HPE and Cisco, through their respective AI Factory, Private Cloud AI and Secure AI Factory platforms, can offer greater control, increased data sovereignty and improved cost predictability for production workloads.

Alongside these options, sovereign cloud and neocloud providers are gaining attention as organisations place greater emphasis on data residency and regulatory requirements. The choice needs to reflect the workload and its constraints, with enough flexibility to accommodate future change.

Turn ambition into a practical production plan

This is the space our data and AI consultancy, Oakland, was built for. We start with structured data discovery ahead of infrastructure decisions, then help shape production AI builds designed for portability across public, private and sovereign platforms.

Customer feedback consistently tells us that organisations want less speculation about AI's potential and more practical guidance from a partner who has solved similar problems elsewhere. That's what Softcat is focused on delivering through 2027: helping customers put the data foundations in place, choose infrastructure that fits and turn promising pilots into working services.



Gary Hawkins
Chief Technologist
Hybrid Platforms

Good vibes only

Software creation has become easier, much easier. But that doesn't make software delivery easier.

What happens when software creation moves beyond the software development team? AI coding assistants, low-code platforms and natural-language tools are making it possible for people with little formal development experience to build applications, automations and agents. The barriers to creating a useful tool are falling. On the whole, that's a good thing.

For organisations with growing development backlogs, it opens up an opportunity. A team can address a small workflow problem close to where it happens, drawing on the knowledge of the people who use that process every day. Ideas can be tested without sending every request back to the dev team.

The first version is just that: the beginning

But creating a tool more easily does not make it easier to run. A personal automation here and a shortcut there need little oversight. But once colleagues rely on an application to complete their work, questions of ownership, security and support become harder to ignore. Who maintains it when a process changes? What happens if its creator leaves? Who knows which tools have become essential?

The risk often appears after the initial success, once the changes start coming in. "Can we just tweak this template?", "Can we change the logic here?" If the organisation cannot see that change in dependency, a convenient, time-saving tool becomes a major headache.

> **"Governance should grow with business dependency."**

Not one-size-fits-all

The answer is to find a route that lets ideas move quickly while applying more scrutiny as their impact grows. A personal productivity tool should be simple to try. But an application used by a team needs a clear owner and a support plan. Software that underpins a critical process needs appropriate engineering, security and governance controls before it reaches production.

This approach gives organisations a way to distinguish experimentation from operational software. It also helps creators understand when to involve development, security and IT teams, and what needs to be in place before more people depend on their work.

Good vibes and invaluable engineers

As more people build, development teams can provide the foundations that help them do it safely: approved platforms, reusable delivery patterns, security controls and architectural guidance. They can focus their expertise where the stakes are highest, while giving business teams a practical way to solve smaller problems.

I'm not talking about a contest between professional and citizen developers here. Far from it. The real opportunity is to connect their strengths. People closest to a problem can explore solutions, while specialist teams help promising ideas become secure, reliable and supportable services. If ownership and oversight grow alongside adoption, faster creation can lead to lasting improvement.



Pat Pendle

Data, Automation and AI Proposition Owner

AI's workspace impact could be right under our noses

How to reduce TCO, improve digital experiences, and accelerate business transformation.

AI's impact is being felt everywhere: in the applications we use, the interactions we have, and the workflows and processes that make things happen. This year, Softcat's customers have told us they have two major technology priorities: AI agents and end-user computing.

So how are those related, and what does the future look like?

> **"EUC isn't really about devices. It's about workstyles, applications and corporate data access."**

"How can I save money and deliver the right devices?"

Organisations tell us that end-user computing is a priority because they want to create a working environment that helps their organisation achieve a reduced TCO whilst also enabling their people's productivity and creativity with incredible digital experiences.

When we understand this and engage with organisations, we encourage them to think more broadly about EUC, beyond just "What device do my users need, and how can I save some money on device spend?" EUC isn't really about devices. The device is a means to an end, and whilst important, should be lower down on your list of priorities after 'workstyles', 'applications', and 'corporate data access'. This ensures that you right-size the technology to your end-users needs, therefore achieving both cost savings and great user experiences.

How does AI fit into the story?

But what about AI? For most information workers today, almost all of their AI usage happens either in the cloud or in the data centre. Very little is processed locally. But cost of AI tokenomics could make the ROI of AI usage hard to justify.

That's where on-device AI and agents come in. It's still early days, but initial testing we've done on workstation-type devices and NVIDIA RTX Spark devices tells us that on-device local AI for real-world use cases isn't far away.

When it does come, maybe in the next year, EUC profiling will need to expand to include the performance required to run AI and agentic workflows locally.

Done correctly, this should deliver on the three key outcomes organisations tell us they want: reduced TCO, improved digital experiences, and business transformation acceleration. Let's get started.



Jack Lewis
Head of Workspace Services

How do we build trust in public sector AI?

The public sector is ready to move beyond AI pilots. The challenge is making adoption work for services, staff and the people who rely on them.

AI is moving up the agenda for public sector leaders. Softcat's latest CES data identifies it as a leading priority among our public sector customers, reflecting both the pace of innovation and the undoubted pressure to achieve more with constrained budgets. The ambition is clear, but turning that ambition into better services takes more than just access to the right technology.

Organisations are starting from different positions

Some are using AI features within off-the-shelf software-as-a-service (SaaS) applications. Others are building capabilities tailored to their organisation through enterprise solutions, or introducing general-purpose AI assistants. Many are combining all three approaches, with different levels of experience and readiness across their teams.

As experiments give way to wider deployment, the focus is shifting to trusted adoption: introducing AI safely and building confidence as its use grows. That means looking beyond individual pilots to the people, processes, policies and technologies needed to support AI across services.

Data is a practical starting point

Our CEX findings also highlight data strategy, integration and service improvement as priorities. Strong foundations for organisational and customer data help teams use AI in ways that support the service they are trying to improve. Those foundations need attention before adoption can scale.

> "Turning that ambition into better services takes more than access to the technology."

Workforce readiness matters just as much. Leaders need to understand the changes AI will bring, while staff need the skills and confidence to work with it. An organisation can introduce new tools quickly, but helping people use them effectively requires sustained attention to change.

As adoption grows, AI needs a clear operating model. Governance, risk management, assurance, financial control and ethical oversight all have a part to play. Service teams must take ownership of the outcomes and decisions involved. AI needs to become an organisational capability, with responsibility extending beyond the technology team.

The benefits are already taking shape in public services. AI-powered contact centres can handle routine enquiries through digital channels around the clock, freeing people to deal with more complex or sensitive cases. The value comes from using that capacity well, so frontline teams can focus their expertise where it matters most.

AI can also help organisations make better use of their data. By supporting more proactive and preventative interventions, it offers a way to respond before demand has fully developed. This connects the investment to a clear service outcome, rather than leaving adoption as an end in itself.

> "AI needs to become an organisational capability, with responsibility extending beyond the technology team."

As we move into 2027, public sector leaders have an opportunity to turn early enthusiasm into lasting improvements. Strong data foundations, a prepared workforce and effective governance give AI adoption a firmer footing. Success will be measured in better services, greater productivity and more time for the work that needs human expertise.



Paul Fleming
Public Sector Advisory Lead

Agentic AI: Financial services' next operating model

While many firms spent 2025 and 2026 piloting generative AI for productivity, operational processes and content, the next phase of transformation is already emerging: agentic AI.

For financial services leaders, agentic AI represents a significant evolution. Unlike traditional AI assistants that generate content or answer questions, AI agents can understand objectives, orchestrate workflows, interact with multiple systems and execute tasks within defined controls. In effect, they become digital colleagues capable of supporting complex and often customer-facing processes.

The opportunity is considerable

Firms face mounting pressure to reduce costs, improve customer outcomes and choice, and respond to increasingly sophisticated fraud and financial crime threats, all whilst maintaining operational resiliency and regulatory compliance.

Agentic AI can help automate activities such as customer service, complaints handling, financial crime investigations, lending operations, regulatory reporting and IT service management. By reducing manual effort and administrative friction, firms can redirect employees towards higher-value activities that require judgement, empathy and expertise. This positions AI as a once-in-a-generation opportunity to fundamentally change both customer and employee experience.

> **"Agentic AI isn't the future, it's happening right now. And it is rapidly becoming the next operating model for the industry."**

How to innovate with AI

Organisations are seeking innovation, but quite rightly only where it is secure, governed and delivers measurable outcomes. Agentic AI comes with a plethora of specific risks that must be managed.

Success will depend on solid foundations. More so than ever, data remains the lifeblood of modern organisations, and secure, trusted information is essential to unlocking business value from AI. Firms must establish robust governance, clear accountability, resilient platforms and strong security controls before deploying agents at scale.

Over the next 12 months, I expect the conversation to shift from "Where can we use AI?" to "Which business outcomes should AI own?" with a healthy dollop of "How do we manage the risks of agentic AI?"

Agentic AI isn't the future, it's happening right now. And it is rapidly becoming the next operating model for the industry. The firms that move first, with the right guardrails, will gain operational efficiency, improve customer experience, and create a lasting competitive advantage.



Katharine Wooller

Financial Services Chief Strategist

Innovation has to earn its keep

Moving from experimentation to execution changes the economics.

As AI becomes embedded in everyday work and production services, the question is no longer simply what it can do, but what value it creates and what it costs to run.

Innovation still matters. But it has to show its value.

Adam Harding
Head of OCTO

Tokenomics and the cost of AI across cloud and platforms

AI token spend is spiralling. But there are ways to move into production at scale and within budget – and with a clear idea of up-front costs.

Every AI conversation has a price tag attached to it, yet pricing remains poorly understood in many AI strategies. Costs are estimated at the point of purchase, then recalculated once the workload goes live. By then, usage may already be growing faster than the budget.

As organisations move from pilots into production, they need a clearer view of what they are paying for, how costs change with demand and when a different deployment model might make sense.

"What's your priority?"

Softcat's CES 2026 survey said...

41% say AI agents

32% say cost control and budgeting

Look beyond the headline token price

Pay-per-token pricing is essentially retail. Convenient, but the costs rack up at scale. And not all tokens carry the same price. **Published rate cards** distinguish between input and output, with generated output often costing several times more per token. The ratio varies by model and service. That price difference does not, on its own, reveal the provider's margin – what matters to the customer is how the rates combine with actual usage to shape the bill.

Plan for the ongoing cost of inference

Inference is the work a trained model does each time it responds to a request. A training job has an end point; inference costs continue to accumulate as people use the service.

In its **October 2025 forecast**, Gartner projected that inference would account for 55% of worldwide AI-optimised infrastructure-as-a-service spending in 2026, rising to more than 65% by 2029.[1] That points to the growing importance of the cost of running AI, well beyond the initial pilot.

Find the right fit for production demand

Hyperscale public cloud AI environments are a great fit for many workloads, but if you're looking for more predictability in your token pricing, other options exist. Dedicated GPU capacity from a specialist neocloud or private infrastructure may offer a lower cost per token for suitable models and workloads. That depends on keeping the capacity productively used.

The comparison needs to include model requirements, request volumes, input and output lengths, performance targets and the full cost of operating the service. Idle GPUs, support effort and software costs can change the calculation. A lower token cost also needs to deliver the quality the workload requires.

> **"Cost per token should be a metric organisations forecast against, not a surprise on the next invoice."**

Softcat's expectation for 2027: cost per token becomes a standard metric organisations track and forecast against, alongside latency and throughput, not an afterthought discovered on next month's invoice. The organisations managing this well are treating inference economics as a discipline in its own right, measuring cost by model and workload before scaling, not after.

Cloud, sovereign or private, the platform is only half the decision. The other half is understanding what every token actually costs, and where. More models and more prompts mean more spend, with no controls. Softcat helps you predict and plan AI costs, keep data secure, and help build proprietary agents that cut your reliance on large, expensive LLMs.



Gary Hawkins
Chief Technologist
Hybrid Platforms

1 Gartner, Gartner Says AI-Optimized IaaS Is Poised to Become the Next Growth Engine for AI Infrastructure, 15 October 2025.

The mic drop moment

What happens when you reduce TCO to reinvest in transformation? And is it really as simple as that?

To the untrained eye, the AI boom might suggest that IT budgets are bigger and better than ever. But rising operating costs, security demands and the pace of change leave many organisations trying to do more with the same investment.

Softcat's Business Tech Report reflects exactly that, with 32% saying cost control and operational efficiency are among the priorities facing UK&I organisations.

Reducing total cost of ownership (TCO) can help. The aim is to release money from technology that no longer earns its place, while protecting the experience and capabilities people rely on. Used well, cost optimisation creates headroom for transformation.

> **"Funding released from unused licences or duplicated platforms could support a Windows 11 transition, modern IT service management or a better employee experience."**

Find the costs hiding in plain sight

Technology estates don't become complex overnight. Devices reach the end of their useful life at different times, applications overlap, and licences continue even after roles have changed. At the same time, ageing infrastructure and fragmented collaboration tools add support effort. Each cost may look modest in isolation, but together they can absorb a chunky share of the budget.

Getting a clear view of the workspace is the starting point:

- Good IT asset management can show what the organisation owns and uses.
- Licence reviews can identify unused or duplicated subscriptions.
- Device lifecycle services can help teams plan replacement and support more consistently.
- Platform consolidation and automation can reduce the effort needed to run the estate.

These measures can have an immediate impact on costs, but they need to be judged by their long-term savings. Removing a tool that employees depend on may simply move the cost elsewhere through slower work or more support requests. The stronger test is whether each change lowers the total cost of delivering a good, secure experience.

Give the savings a purpose

Savings have the greatest impact when there is a plan for reinvestment. Funding released from unused licences or duplicated platforms could support a Windows 11 transition, modern IT service management or a better employee experience. It might also help an organisation test AI services against practical use cases, with the right governance in place.

Leaders need three things: a baseline for current costs, a way to measure the effect of changes and a clear link between savings and the next investment. That makes it easier to show whether optimisation is actually improving outcomes rather than simply reducing a budget line.



Kelly Calver
Head of Workspace



Jack Lewis
Head of Workspace Services

Build a cycle of improvement

Modern tools and simpler processes can lower ongoing support costs, but that's just the first step. Going further means creating better experiences that can improve adoption, helping organisations realise more value from what they buy. The organisations seeing the greatest success recognise that cost optimisation and transformation are not competing priorities. When approached strategically, they reinforce one another. This is the mic drop moment.

The goal is a workspace that costs less to run and does more for the organisation. By finding waste, protecting what works and reinvesting with intent, businesses can turn TCO reduction into a practical route to long-term change.

Understand before you accelerate

Finding the money to transform is only part of the challenge.

Years of accumulated technology, fragmented data and interconnected systems can make adding something new easier than understanding what is already there.

Before accelerating, organisations need to understand the foundations beneath them, simplify where they can and know how data, decisions and work move across the business.

Adam Harding
Head of OCTO

Talk data to me

Conversational analytics is shortening the distance between a business question and an answer.

Businesses have spent years being told to listen to their data. There was just one small problem: the data wasn't particularly good at talking back.

Ask a seemingly simple question in a meeting and you'll probably recognise what happens next. Why did sales drop last month? Someone opens a dashboard. It doesn't quite answer it. Another view is needed. Perhaps another data source. The request goes to BI, joins the queue and, eventually, a new report appears. By which point, you've got another question.

That gap between having a question and getting an answer is now collapsing remarkably quickly.

Your data can finally talk back

Modern data platforms are changing what businesses can expect from their data. Microsoft is building these capabilities into Fabric, Databricks has Genie, Snowflake has Cortex, and Google is bringing Gemini-powered conversational analytics into BigQuery.

Different platforms, but all making it possible for people to ask questions of their data directly. Why did sales fall last month? Which region drove that? Was it one product line? Show me the trend for the last 12 months.

We've been here before. Sort of.

Putting data in one place isn't new. What is changing is that modern platforms increasingly combine data storage and processing with governance, intelligence and AI.

But there is some less glamorous work involved first.

"Data is the foundation of AI" may be one of the oldest lines in the industry. In 2026, it is still true. We use a three-stage framework to get data ready for AI analytics: trust the critical data, connect it through a modern platform, then use conversational analytics to turn it into valuable answers.

Trust it. Connect it. Then get value from it. Together, these principles mean focusing governance and engineering on the data products, workflows and decisions where better answers will make a measurable difference.

What is an answer worth?

There is another side to this conversation: cost. Modern data platforms consume compute, and those costs can fluctuate. Softcat helps customers secure competitive pricing and funding through its scale and relationships with major platform vendors, while Oakland, our specialist Data & AI arm, designs and delivers the transformation. Together, we connect the business case, platform choice and commercials with the essential work on the data underneath, then build something people will actually use.

> **"Trust it. Connect it. Then get value from it."**

Conversational analytics is only going to get better from here. The organisations that benefit most won't necessarily be the ones chasing every new feature. They'll be the ones getting their data ready to take advantage of them.

Your data probably already has plenty to say. It might be time to start talking to it.



Andrew Hermsen

Chief Technologist Data, Automation and AI

Following the breadcrumb trail to answer: Where is the work?

How the trail of data, decisions and value in a world of AI is completely reshaping what work is.

We have become pretty good at understanding where our technology lives and who uses it. We can see applications, servers, devices, networks and users. But as AI becomes embedded in everyday operations, a more fundamental question is emerging: where is the work?

Looking back over previous years, most work was performed by people, and those people became increasingly supported by technology. Today, that is starting to flip round: work is being done by technology, which is in turn being supported by people – quite a change!

Following the trail of data, decisions and value

Work increasingly flows between people, applications, supply chains, automation platforms, AI models and agents. A customer request might trigger a workflow, move data between systems, call on several AI services and lead to decisions made with different levels of autonomy.

This is why orchestration and integration have so rapidly become top priorities. The challenge is no longer getting the work done. It is understanding how that work moves across an increasingly complex terrain of people, processes and systems, and asking:

- How does data move?
- Where are decisions being made?
- Which workflows cross organisational boundaries?
- What consumed the budget?
- How do I know when something has gone wrong?
- Who is accountable when something goes wrong?
- What is my plan B now that people don't do this job anymore?

Keep the context with the activity

Answering those questions takes more than a view of whether individual systems are running. We need to preserve context as work moves: what triggered an action, which information informed a decision and what happened next.

Metadata starts the breadcrumb trail, and integration allows it to cross boundaries. Orchestration coordinates the activities and hand-offs that follow. Without these things, we may be able to see the woods, but miss the trees.

> **"We need to preserve context as work moves. Organisations that can follow the trail of data, decisions and value will be better placed."**

Lost in the woods?

As AI becomes more deeply embedded, it will be critical to understand the relationships between our people, processes and technologies. The ultimate goal? To better understand and audit, how and where work is being done.

Organisations that can follow the trail of data, decisions and value will be able to maintain accountability, efficiency and control. Those that can't, will get lost in the forest.



Andrew Hermsen

Chief Technologist
Data, Automation and AI

Networking at machine speed

How AI is reshaping network traffic – and the actions you can take right now.

AI is changing the nature of network traffic. Look no further than Cloudflare's numbers, which suggest that the tipping point has already been reached; bot traffic outnumbers human traffic.

In **Cisco's analysis**, an agent completing a task generated up to 450% more total traffic than a person performing the same task manually. Separately, Cisco forecasts that AI inference will account for 25% of total network traffic by 2035.[2]

It's not a surprise. Agents can call tools, move data and query models repeatedly, creating demand that extends beyond individual human interactions. Networks need to support these changing patterns alongside higher capacity, security and consistent policy enforcement.

Volume is only part of the challenge. AI can create sustained connections and repeated exchanges between applications, models and data sources, operating continuously at machine speed and creating persistent network demand.

Start now by mapping where AI applications send and retrieve data. Measure traffic volumes, connection duration, latency and upstream demand, then check whether critical paths have the capacity, resilience and security controls they need.

[2] AI Impact on Wide Area Networks: Cisco 2026 report
[3] Alert Fatigue Reduction with AI Agents | IBM
[4] AI Cluster Cost Breakdown: CapEx (2026) | American Compute

Use agentic NetOps with clear boundaries

AI is also changing how networks are managed. Agentic network operations, or NetOps, extends AI's role from analysing conditions and recommending responses to taking action within defined limits and with human oversight. Citing Vectra AI research, IBM reported that SOC teams handle an average of **4,484 alerts per day, with 67% ignored due to false positives and alert fatigue**.[3]

Agents can monitor network health, identify anomalies and recommend remediation. Where appropriate controls are in place, they can also carry out approved corrective actions, helping teams respond faster while reducing routine, repetitive work.

> "Measure how traffic is changing, introduce automation with clear controls and assess network investment by the performance it enables."

Let's take a common use case, such as investigating recurring alerts. First, agree which actions require approval, record what the agent does and establish a rollback process before extending its authority.

Design the datacentre network around the workload

Within AI datacentres, network performance can affect the value of the whole platform. American Compute estimates that networking accounts for around 20% of AI cluster hardware costs, based on the bills of materials it has reviewed.[4] If data cannot move quickly enough between compute resources, job completion times can rise and GPUs can spend more time waiting.

Distributed AI workloads need reliable, low-latency connections and effective congestion management for east-west traffic between systems. Ethernet using RoCEv2 and InfiniBand are options for these environments, while Ultra Ethernet specifications target the demands of AI and high-performance computing.

Before committing to an architecture, test representative workloads and compare job completion times, GPU utilisation and total operating spend. Include the skills, support and hardware availability needed to run it.

Being prepared for AI means understanding every workload. Measure how traffic is changing, introduce automation with clear controls and assess network investment by the performance it enables across the AI platform.



Thomas Rowley
Chief Technologist
Networking & Connectivity

Trust must be designed in

Understanding how technology, data and work connect is essential, but it does not make them trustworthy.

As AI takes on more activity and decisions move between people, agents and systems, questions of identity, security, sovereignty, accountability and resilience become harder to separate.

Trust can no longer be added at the end. It has to be designed in from the start.

Adam Harding
Head of OCTO

How data and AI sovereignty will shape your strategy

Where your AI runs, who can access its data, and who controls the service are becoming board-level decisions.

Data location has long been a minor checkbox on a compliance questionnaire. And while GDPR moved the dial on consumer data location, it will be the next 12 months that sees the biggest shifts. In 2027 and beyond, we expect data sovereignty to play a much bigger role in business strategy.

Tightening data residency rules, the EU AI Act's transparency obligations, and rising geopolitical scrutiny mean organisations can no longer treat "the cloud" as one undifferentiated place to put data and workloads. Where AI processes and stores information, and who can access it, is now central to risk management, not an afterthought to it.

Look beyond where the data sits

Data residency is one part of that picture. A workload may run in a local data centre while support arrangements, administrative access or operational dependencies extend beyond that location. Understanding sovereignty means examining those relationships as well as the address of the server.

AI governance adds further responsibilities. The EU AI Act, for example, introduces **transparency obligations** for certain AI systems. Those obligations should be assessed alongside the separate legal, contractual and sector requirements governing where data can be held, processed and accessed.

Match the platform to the level of control required

This creates a choice of approaches. Hyperscalers are introducing sovereign cloud offerings with specific governance and operational controls. **AWS's European Sovereign Cloud**, for example, uses dedicated European legal entities and infrastructure that is physically and logically separate from its other regions.

Specialist sovereign providers and some neocloud providers offer further options. Private or hybrid infrastructure can also give organisations more direct control over sensitive workloads. In every case, the detail matters: the provider's location or a sovereign label alone cannot establish whether a service meets a particular requirement.

Start with the data before committing to a platform

Softcat's prediction for 2027 is that sovereignty requirements will increasingly shape infrastructure choices from the start. Treating them as a later configuration exercise risks rework, delays and gaps in compliance.

The practical starting point is the same one that matters for any AI production build: before committing to a platform, understand what data you hold, how sensitive it is, and where it legally needs to live. That classification work is what our data and AI consultancy, Oakland, focuses on every day.

Oakland's data governance specialists bring this to every engagement, so infrastructure decisions are made with sovereignty from the ground up, not tacked on retrospectively.

Getting this right isn't just about avoiding penalties. It's what lets organisations deploy AI with genuine confidence, knowing their data and their customers' trust stay exactly where they should.

> "Sovereignty requirements should shape infrastructure choices from the start."



Gary Hawkins

Chief Technologist
Hybrid Platforms

Can you trust AI?

Let's talk about AI governance. How to build security, assurance and resilience into your AI projects.

AI is rapidly moving from innovation projects into day-to-day business operations. As organisations look to unlock greater productivity, automation and insight, the conversation is shifting from "How quickly can we adopt AI?" to "How confidently can we trust it?"

Organisations must focus on three critical areas in order to deliver safety and governance around any AI deployment: security, assurance and resilience.

> **"Trust will be one of the biggest differentiators between organisations that successfully scale AI and those that struggle to realise its value."**

Why this matters now

Many organisations are already using AI more extensively than they realise. Employees are experimenting with AI tools, technology providers are embedding AI into existing platforms and business leaders are looking for ways to drive efficiency and competitive advantage.

As adoption accelerates, so do the challenges. Organisations must be confident that AI systems are secure, that outcomes can be trusted, and that critical business processes can continue if AI services fail or produce unexpected results.

Trust will be one of the biggest differentiators between organisations that successfully scale AI and those that struggle to realise its value.

Understanding the challenge

- AI security focuses on protecting AI systems, data and models from misuse, compromise and emerging threats.
- AI assurance provides confidence that AI is operating as intended, with appropriate governance, transparency and accountability.
- AI resilience helps ensure organisations can continue operating effectively when AI systems are unavailable, disrupted or produce inaccurate outputs.

Softcat's view of the future

Over the next few years, we expect AI security, assurance and resilience to become embedded within wider business operations, much like cyber security and operational resilience are today.

Leading organisations are already moving towards continuous AI governance, greater board-level oversight of AI risk, and resilience models that account for AI-powered business processes and third-party AI dependencies.

The organisations that gain the greatest value from AI are unlikely to be those that adopt it the fastest. They will be those that adopt it most responsibly.

By focusing on security, assurance and resilience from the outset, businesses can create the trust required to scale AI safely, unlock its full potential and build a foundation for long-term transformation.

How to gain visibility and control

Key priorities should include:

- Understanding where AI is already being used across the organisation.
- Establishing clear governance and ownership.
- Defining policies for the safe use of AI and business data.
- Extending cyber risk management processes to include AI-related risks.
- Educating employees on the opportunities and responsibilities associated with AI.



Trevor Highley

Head of Networking and Cyber

How do you take control of AI agent permissions?

AI identity, governance and access management is essential for giving every agent an owner, appropriate permissions and a clear record of its activity.

Who, or what, is accessing your services and data? Security teams have traditionally focused on human identities and service accounts. AI agents add another type of identity to govern. They can access data and take action across business systems, making control over their permissions essential.

Discover agents and limit their access

As agents become more autonomous, identity provides an important point of control. Security teams need to discover AI agents, enforce least-privilege access and review permissions throughout each agent's lifecycle.

Governance platforms are evolving to manage agents alongside workforce identities, using access reviews, approval workflows and conditional access policies where supported.

Beyond authentication

Authentication is only part of the challenge. AI agents interact with APIs, large language models (LLMs) and Model Context Protocol (MCP) integrations.

Access decisions need to reflect the risk, context and business purpose, while limiting excessive permissions, unauthorised behaviour and privilege escalation. Just-in-time access, ongoing validation and behavioural analysis can all play a part here.

Make ownership and oversight visible

An AI identity framework needs to bring governance, security and compliance together. That means maintaining a central inventory of agents, linking each to a responsible owner and recording activity in comprehensive audit trails. Frameworks such as the **NIST AI Risk Management Framework and ISO/IEC 42001** can help inform the wider governance approach.

Technology leaders need to enable innovation while keeping AI systems within defined security, regulatory and ethical boundaries. Clear ownership and appropriate access give teams a practical foundation for doing both.

> **"Every AI agent needs a responsible owner."**



Ryan Birch
Cyber Tower
Sales Lead

Secure networking needs to be built in, not bolted on

Consistent policy and enforcement start with security baked into the network itself.

You probably have no shortage of security controls. Over time, you have invested in firewalls, access policies, monitoring platforms and cloud security services. Yet despite that investment, applying policy and maintaining visibility across the entire environment can still be difficult.

The roller coaster never stops: right now, AI is introducing new traffic patterns, access requirements, and data flows. Each decision to add another security tool may have made sense at the time, but together they can leave a fragmented stack that lacks control and visibility.

Close the gaps between controls

This is why secure networking is becoming a more important conversation. Many organisations are reaching the limits of managing disconnected controls across offices, cloud platforms and remote users. The challenge is to make those controls work consistently wherever people, applications and data connect.

A policy that is clear in one part of the estate may be difficult to enforce in another. Separate monitoring tools can provide different views of the same activity, while unclear ownership makes it harder to decide who should respond. Adding capabilities without resolving those gaps can increase the effort needed to manage them.

Make security part of the architecture

Security needs to be considered when architectural decisions are made. The network has an important role in providing visibility, enforcing access policies and separating workloads through segmentation. Those capabilities need to work with identity systems and the wider security environment.

That means bringing identity, policy, visibility and enforcement into the same design conversation. For example, decisions about a new cloud application should include who can access it, how that access is controlled and how activity will be monitored. Addressing these questions together makes consistent operation easier as the estate grows.

Reduce the complexity teams have to deal with

The operational pressure is already making itself known: Softcat's Networking & Connectivity research found that 70% of NetOps teams are experiencing alert fatigue.

When visibility is fragmented and alerts lack context, teams often spend more time interpreting information and coordinating responses instead of taking immediate action. Better integration and clearer ownership can help them act with greater confidence.

The organisations making the most progress are not necessarily investing in the most controls, they are instead reducing fragmentation. They're making policy more consistent, visibility more comprehensive and enforcement easier to manage across the whole environment.

Secure networking is ultimately a response to that challenge. It reflects a growing recognition that consistency, not capability, is becoming the defining factor in security and operational resilience.

> **"The challenge is to make those controls work consistently wherever people, applications and data connect."**



Ebow Barnes

Networking & Connectivity Sales Manager

Post-quantum cryptography: Why 2027 is the year to get started

The hardest bit isn't picking the new algorithms; it's knowing where your cryptography actually lives.

Quantum computing will eventually break the public-key encryption that sits behind pretty much everything we all do online: payments, customer data, digital signatures, remote access and much more.

That day might still be years away, but the risk is already here. Attackers are harvesting encrypted data now so they can decrypt it later, so anything that needs to stay confidential for the next ten years is already exposed.

Post-quantum cryptography (PQC) readiness sits right where those two meet. Yet it's rarely on the agenda.

The direction from regulators

The direction from regulators is clear. The National Cyber Security Centre (NCSC) expects organisations to have identified the cryptography that needs upgrading and to have a migration plan by 2028, with high-priority systems moved by 2031 and full migration by 2035.

Looking at 2027, I think we will see two big shifts. First, PQC moves out of the security team and into the boardroom, as regulators and auditors start asking to see roadmaps.

Second, certificate management becomes the first real battleground. Public TLS certificate lifetimes drop to 100 days in March 2027 and 47 days by 2029, so organisations will be forced to automate – which is a good first step for quantum readiness anyway.

Start with discovery

The hardest bit isn't picking the new algorithms; it's knowing where your cryptography actually lives. Start with discovery and build an inventory across your applications, infrastructure and third-party services. Then prioritise systems by how sensitive the data is and how long it needs protecting, ask key suppliers about their PQC roadmaps, and pilot hybrid approaches on a low-risk system first.

Most importantly, design for agility: centralised key management and automated certificate rotation mean the next change is routine rather than a full re-architecture.

Migrate on your own terms

The organisations that start in 2027 will get to migrate on their own terms. The ones that wait could end up paying a much greater price.

> "Attackers are harvesting encrypted data now so they can decrypt it later."



Trevor Highley

Head of Networking and Cyber

Keeping critical public services running

Resilience in the public sector starts with knowing which services need to survive the most. But who makes that decision?

Resilience has become a critical priority for public sector organisations. Geopolitical instability, increasingly capable cyber threats and the rapid adoption of AI are adding to the pressures facing essential services. Softcat's latest CEX data reflects that shift, with resilience becoming a growing concern among our public sector customers.

The Cyber Security and Resilience (Network and Information Systems) Bill adds to the focus on protecting essential services. Its proposals underline the importance of stronger cyber defences, but organisations already have a practical reason to act: people depend on their services being available when they need them.

Boards and audit committees are increasingly asking for assurance and evidence. Technology leaders need to explain what disruption would mean for services, which capabilities are essential and where weaknesses remain. Those answers make resilience a discussion about organisational priorities and public trust.

Traditional business continuity and disaster recovery approaches are evolving to include the idea of a Minimum Viable Organisation. This means defining the essential services and capabilities that must remain available during significant disruption. It gives leaders a practical starting point for deciding what the organisation must protect and restore.

> "Resilience starts with knowing which services cannot stop."

That language matters. A Minimum Viable Organisation is a concept that chief executives and service leaders can engage with directly. Defining what must keep running brings resilience into decisions about service delivery and organisational responsibility. Resilience starts with knowing which services cannot stop.

Recovery must also be controlled and trusted. Restoring availability is only part of the task: leaders need confidence in the systems, services and data they bring back into use. That makes the quality of recovery as important as the ability to recover at all.

For technology teams, the work spans security operations, threat detection, incident response, reporting and incident management. The National Cyber Security Centre's Cyber Assessment Framework provides a structured way to assess how cyber risks to essential functions are being managed. Its supply chain principles also bring attention to the organisations and services on which delivery depends.

That requires a clear understanding of the application landscape, whether services run through software-as-a-service (SaaS), cloud platforms or hybrid environments. Leaders also need assurance about the infrastructure underneath. Understanding gaps, weaknesses and dependencies makes it possible to judge their effect on critical services. A weakness only becomes meaningful to the wider organisation when its impact on service delivery is understood.

> "A weakness only becomes meaningful to the wider organisation when its impact on service delivery is understood."

Treating resilience as a strategic capability gives these activities a shared purpose. The outcome is continued access to critical public services, confidence in the organisation and protection for the integrity of people's data. Keeping that purpose in view helps leaders make better decisions before disruption arrives.



Paul Fleming
Public Sector Advisory Lead

Adoption is the differentiator

Better technology, stronger foundations and more sophisticated controls create the right conditions for transformation. But none of them determine what an organisation chooses to do with that transformation.

Someone still has to decide what should change, what should be automated, where decisions belong and where human judgement matters most. Ultimately, the question isn't simply how technology will change. It's how work should change with it.

Adam Harding
Head of OCTO

AI agents and human expertise: Better together

Why the next advantage comes from collaboration, not replacement.

Whenever a new technology emerges, the conversation tends to begin with a binary question.

Will it replace people?

We saw it with automation. We saw it with low-code platforms. And now we're seeing it with AI agents. Yet the organisations creating the most value from AI today are asking one vital question:
How can people and agents produce better outcomes together than either could achieve alone?

AI is evolving from tool to teammate

AI agents are not like standard AI chatbots. Rather than responding to a single request, they can increasingly plan, coordinate, retrieve information, interact with systems and complete sequences of actions with limited human intervention.

Much of the hype surrounding AI focuses on machines replacing humans. In reality, value often comes from recognising that humans and agents excel at different things.

AI agents typically bring:

- Speed
- Scale
- Consistency
- Pattern identification
- Continuous availability
- Execution of repeatable activities

People continue to bring:

- Judgement
- Context
- Empathy
- Creativity
- Accountability
- Relationship building

Neither is inherently better. They're simply different.

The organisations seeing the greatest benefit from AI are learning to design work around these complementary strengths instead of treating AI as either a replacement or an assistant.

Better outcomes through collaboration

The most interesting AI stories aren't about automation. They're about collaboration.
Let's take an example from customer service in the insurance sector. An AI agent can instantly pull together customer history, identify relevant policies, summarise previous interactions and recommend potential resolutions.

But when a customer is frustrated, when a conversation becomes sensitive, or when trust needs to be rebuilt, human judgement matters. The best outcome often comes from combining machine efficiency with human empathy.

Some questions will become increasingly important:

- Which activities should be handled by agents?
- Which decisions require human oversight?
- Where does accountability remain?
- How should success be measured?
- What skills will people need to collaborate effectively with agents?

Building the workforce of the future

In most cases, agentic AI increases the need for human skills.

Here's why: as routine and repetitive work becomes easier to automate, the value of judgement, creativity, leadership and specialist expertise becomes more visible.

Skills such as critical thinking, communication, collaboration and decision-making are likely to become even more important because they help people govern, direct and enhance the work that agents perform.

Perhaps the biggest shift leaders need to make is moving beyond the replacement debate altogether. The future is unlikely to belong to organisations that rely entirely on people. Nor is it likely to belong to those that attempt to automate everything.

It will belong to organisations that understand how to combine human judgement with machine speed, human creativity with machine scale, and human accountability with machine execution.

> **"The best outcome often comes from combining machine efficiency with human empathy."**

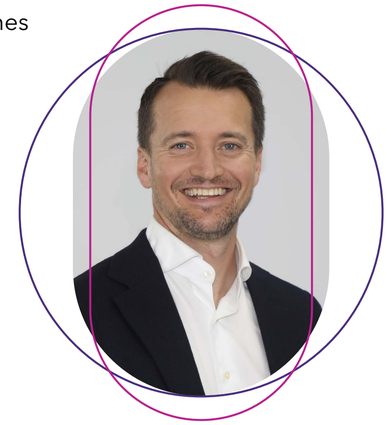

Adam Fraser
Corporate Sales Director

How should work happen today?

We will not compete on intelligence. We will compete on execution.

Much of the AI conversation has focused on capability. Organisations have invested in copilots, assistants, agents, and increasingly advanced models. But as those tools become more accessible, I believe the competitive advantage will come from how effectively organisations put them to work. Access to intelligence alone will not set a business apart. Execution will.

The great misunderstanding of AI

It's a simple equation: knowledge worker + AI assistant = greater productivity.

This isn't wrong, but it is incomplete. It is local optimisation. The worker may become more efficient, but the system (or department) does not. Organisations are no longer constrained by thinking or intelligence; they are constrained by what they choose to do with it.

The individual may become more efficient while the wider team or department continues to work in much the same way.

If someone can produce an analysis in minutes but the decision it supports still waits a week for approval, the organisation has only addressed part of the problem. The next question is what needs to change that person for the whole process to benefit.

Redesign how work happens

The bigger opportunity is to rethink work in light of what AI makes possible. Alongside asking where AI can improve an existing task, we should ask: what can we do today that we could not do yesterday?

That shifts the focus towards systems of work. Some decisions can be supported by clearer criteria and better information. Repetitive activities can be automated, and agents can act within defined boundaries. People can give more attention to judgement, creativity, and the exceptions that need their expertise.

Operating models then become as important as technology platforms. Governance, accountability, and economics move to the foreground. Who owns the process and its outcome? Where should decisions sit? What authority should be delegated? How will we measure the value created?

> **"What can we do today that we could not do yesterday?"**

Bring people into the change

We are still working out how to prepare our people for this shift. As AI capabilities evolve and automation becomes more commonplace, there is a difficult balance to strike. How do we help subject matter experts feel empowered to drive change today while remaining confident in their long-term value?

Their knowledge is essential to answering the question. They understand how work actually happens, where the delays sit and which exceptions demand care. Involving them in redesigning that work gives us a better starting point than making assumptions about what should be automated.

I certainly don't have all the answers. But I am certain of one thing: we will not get far without our people on board. We need to understand how work was done yesterday so we can ask how it should happen today, and that is how we prepare for the bigger question: what will work look like tomorrow?



Andrew Hermsen

Chief Technologist Data, Automation and AI

The future of commercial operations

Procurement is evolving from a cost-control function into a strategic enabler of business transformation.

Our Customer Experience Survey highlights that while price competitiveness (85%), customer service excellence (78%) and value for money (70%) remain the leading priorities for procurement teams, organisations are increasingly looking beyond cost alone. Trusted expertise, delivery performance, governance, innovation and social value are all shaping buying decisions.

At the same time, procurement teams are being asked to support many of the same priorities explored throughout this report: accelerating AI adoption, controlling costs, improving resilience, reducing complexity and delivering measurable business outcomes. The challenge is balancing commercial efficiency with long-term strategic value.

In this section, we explore how procurement is responding to that shift. From navigating global supply chains and governance requirements, to embedding social value into procurement decisions, and understanding the role AI will play in the future of bidding and supplier engagement, procurement has become a critical driver of business success.

The organisations creating the greatest value are no longer asking, "How do we spend less?" They're asking, "How do we get more from every investment?"

Mark Brazington
Sales Operations Director

Procurement contacts

85% Price competitiveness

78% Customer service excellence

70% Value for money

64% Quality of advice and expertise

64% Delivery times and turnaround

Global standards meet local realities: Rethinking multinational tech procurement

Creating a global technology strategy is one thing. Making it work consistently in every country is another.

For multinational organisations, technology procurement goes far beyond negotiating global agreements or standardising devices. The new imperative is to establish effective governance: deciding what should be consistent worldwide, where local flexibility is essential, and how to maintain oversight across an increasingly complex technology estate.

Balancing global standards and local flexibility

The answer isn't necessarily a push towards centralisation. It's finding the right balance between global standards and local flexibility. Common governance, technology policies and security requirements provide consistency, while regional sourcing capabilities, local expertise and connected data help organisations navigate individual market requirements.

> "A global technology strategy is only as effective as an organisation's ability to execute it locally."

Multiple threats: Data and supply chains

As businesses look to adopt AI across international operations, new challenges emerge. Who can access your data? How do different permissions either enable or restrict employees in their daily tasks? And what does your 'data supply chain' look like?

Geopolitics adds further uncertainty. Changing tariffs, trade restrictions and disrupted supply chains can alter the cost, availability and delivery of technology with little warning. Organisations cannot predict every disruption, but they can prepare to respond.

How to measure procurement's effectiveness

Looking towards 2027, I believe successful multinational organisations will increasingly measure procurement effectiveness through more than cost savings. Visibility, resilience, regulatory compliance and the ability to deploy technology consistently across borders will become equally important.

Ultimately, a global technology strategy is only as effective as an organisation's ability to execute it locally.

Softcat's 2026 Customer Experience Survey reflects these pressures:

- **52%** say cost control and budgeting is a priority
- **42%** say IT standardisation and rationalisation is their priority
- **37%** say governance, risk and compliance is their priority



Liam Dennison
Head of International

How MAT is reshaping public sector procurement

Social value needs to shape the bid from the outset and stand up during delivery.

The Procurement Act's move from "most economically advantageous tender" (MEAT) to "most advantageous tender" (MAT) reinforces that value can extend beyond price and technical quality. Alongside this, social value is becoming more closely connected to business objectives.

For the technology sector, this is an exciting opportunity. Every day, we help organisations use technology to transform the way they work, but we can also use that same technology, alongside the strength of our supply chains and partnerships, to help tackle challenges around digital inclusion, employability and economic opportunity.

At a time when so much of everyday life relies on digital access and skills, technology has an important role to play in helping people and communities thrive.

Social value is moving into the foundations

For in-scope central government procurements commenced on or after 1 January 2027, PPN 026 sets a minimum 20% social value weighting for contracts valued at £5m or more, doubling the previous minimum. For specified absolute evaluation methods, this applies to the non-price criteria.

At that level, it cannot be left until the end.

Two suppliers may offer similar technology, pricing and service. Increasingly, the differentiator between them may be the wider impact a supplier can create through jobs, skills, development, apprenticeships, community investment and resilient UK supply chains.

That means we need to think about social value alongside the solution and commercials from the outset, so the commitments are built into the bid rather than added afterwards.

For Softcat, that broadens the story we need to tell.

The value is not just in the hardware, software, or services being bought. It is also in the organisation behind it, the people it supports, the opportunities it creates, and the impact it delivers through its supply chain, partners, and communities.

> "What looks good on paper also has to work in practice."

What happens after the tender is won?

The Procurement Act creates more visibility around supplier and contract performance, while PPN 026 requires social value commitments to be monitored as part of the contract.

That means what looks good on paper also has to work in practice.

Strong social value can help differentiate a bid, but only if we know we can actually deliver what we are promising. Taken together, these changes mean social value is becoming a core part of how we shape public sector bids. For us, the opportunity is to bring social value in earlier, build stronger and more credible commitments around it, and make sure what we promise at bid stage still stands up once the contract is live.

But this isn't about writing better commitments. It's about building a business where social value is already embedded in how you operate, so that when opportunities arise, your bid simply reflects what you do every day.



Dani King
Social Value Lead

How AI is changing bids

When polished answers become easier to produce, evidence and judgement matter more.

I recently reviewed a bid and spotted the word "mangement" instead of "management". And weirdly, I found it quite reassuring. It felt human.

We are now in a world where AI can produce a polished, professional-looking response in seconds. The grammar is correct, the structure is neat and the tone is confident.

But if everyone can produce something polished, the real question becomes: what is actually good?

For me, that is one of the biggest ways AI is changing bidding. It is making the surface easier, which means the substance underneath matters more.

> "But if everyone can produce something polished, the real question becomes: what is actually good?"

Better AI starts with better information

At Softcat, we are not just asking how AI can help us draft faster. We are working on what sits underneath it.

We are exploring how to structure and simplify our bid library so that it is easier for AI to use effectively, with clearer facts, evidence, customer outcomes, credentials and proof points. Alongside this, we are developing AI agents to help assess risk, capture customer feedback at scale, identify recurring themes and feed those insights back into the content library, helping it become more useful and relevant over time.

The aim is not simply to write faster. It is to make the whole bid process learn.

AI can sound convincing, even when it's wrong

AI carries risks, from hallucinations to perfectly decent content becoming more generic, repetitive and corporate.

In Poland, road contractor Exdrog was excluded from a public procurement process after submitting unverified AI-generated information that included references to tax rulings that did not exist. The contract was worth around €4 million. That's a very costly hallucination.

Customers are using AI too

Suppliers are not the only ones experimenting with AI. We are already seeing RFPs and presentation agendas that appear to have been heavily created by AI.

We also need to be mindful that procurement teams may increasingly use AI in evaluation too.

If AI is helping shape the RFP, suppliers are using it to respond, and buyers are using it to assess answers, genuine differentiation becomes even more important.

That makes clear evidence, customer relevance and a strong point of view much harder to fake, and much more valuable.

So what happens to the Bid team?

AI should make parts of bidding quicker. But if the time saved simply allows us to fit in another bid, we have only made the process faster but not necessarily better.

That time can instead be invested in the areas where a good Bid Manager adds the most value: qualification, customer strategy, solution challenge, evidence, win themes, SME engagement and stronger reviews.

AI should make Bid teams less administrative, not less important. Our job is to make sure the bid is strategically sound, commercially sensible, low enough risk to stand behind, compelling to the customer and still recognisably human. Perhaps that's why "mangement" made me smile.



Bonnie Miller

Enterprise Bid Lead

Summary

Turn AI ambition into business value.

As organisations move from experimentation to execution, success will depend on balancing innovation with discipline: deploying AI responsibly and at scale, simplifying increasingly complex IT estates, staying on top of the Cyber threats, and controlling costs while continuing to invest in transformation. All of those together will not be easy to navigate.

Key priorities for 2027

- Scale responsible AI from pilot to production to deliver measurable productivity, customer and operational outcomes.
- Standardise and simplify technology estates to reduce complexity, improve resilience and accelerate change.
- Optimise cost and consumption through better visibility, governance and investment planning.
- Strengthen cyber, data and AI governance to build trust and enable responsible innovation.
- Use technology as a driver of business growth, agility and competitive advantage.

The three actions that matter most

1. **Execute AI responsibly and where applicable at scale**
   Move beyond experimentation and embed AI into business processes, decision-making and customer experiences.
2. **Simplify to accelerate**
   Standardise platforms, reduce technical debt and create a more agile, manageable and secure IT environment.
3. **Optimise to invest**
   Control costs, improve efficiency and unlock funding for future innovation and growth.

> **“The world is not slowing down, but technology and AI should only travel at a speed of responsibility, whilst balancing innovation and accountability. The organisations that lead in 2027 won’t be those adopting the most technology. They’ll be those that successfully balance AI innovations, cyber impacts, IT simplification and cost optimisation to create sustainable business value.”**



Dean Gardner
Technology Director

softcat

01628 403403
Find out more www.softcat.com