FEBRUARY 2026
networkingplus.co.uk

NETWORKING+

Ransomware attacks

Building a human firewall to defend against the rising threat

Irvin Shillingford, Hornetsecurity, p6

Building early warning alerts

Reduce knock-on, escalating disruption

Martin Smillie, SITA, p8

AI in cybersecurity

An opportunity or a risk for MSSPs?

Haris Pylarinos, Hack The Box, p14

Beyond the edge: Orchestrating the UK's software-defined future from 5G to 6G

While the UK's data centres provide the "brain" of the digital economy, the nervous system—composed of WAN, LAN, and the emerging 5G Standalone (5G SA) underlay—is undergoing a simultaneous and radical transformation. The convergence of software-defined architectures and advanced wireless is moving the UK from a "connectivity first" model to one of "intelligent orchestration."

The UK's wide area networking (WAN) landscape hit a critical inflection point this month with the £2 billion acquisition of Netomnia by nexfibre. This merger has created a unified metropolitan fibre giant, offering a direct challenge to BT Openreach's wholesale dominance. For enterprise networking teams, this consolidation simplifies the "underlay" puzzle, providing a single, high-capacity fibre footprint across the UK's major cities to support SD-WAN deployments.

Technical audits from February 2026 show that 100% of the UK's leading SD-WAN vendors have now verified local Point of Presence (PoP) ownership in London and Manchester. This ensures sub-10ms latency for UK traffic—a technical prerequisite for the "Agentic AI" applications currently being trialled in the finance and legal sectors.

This month, the UK reached a significant wireless milestone. Virgin Media O2 (VMO2) announced that its next-generation 5G Standalone network is now live in over 500 locations, including a massive new rollout across Greater Manchester.

Unlike early 5G, which relied on 4G cores, 5G SA is "cloud-native." This shift is critical for SD-LAN integration. For the first time, UK businesses can treat a 5G slice as a seamless extension of their office LAN. This "network slicing" allows IT managers to guarantee bandwidth for critical industrial robotics or sensitive medical data, effectively turning the metropolitan 5G network into a private, software-defined extension of the campus.

While 5G SA matures, the "private network" market has gone mainstream. Currently, the UK leads Europe with over 6,500 private cellular deployments. These installations are moving beyond simple mobile broadband to provide the foundation for AI-Native LANs, where the network autonomously manages frequency congestion and security protocols.

Looking further ahead, the University of Bradford and Digital Catapult secured a combined £13.1 million in new 6G research funding this month. The focus is on the transition from connectivity to sensing. 6G, currently entering the standards-definition phase (3GPP Release 20), is being designed as an "Integrated Sensing and Communication" (ISAC) framework. In this future, the network underlay won't just move data; it will perceive motion and environmental context, turning every metropolitan LAN into a real-time digital twin of the city.

The government's Mobile Market Review reinforces the 2030 ambition: 5G SA coverage for all populated areas. However, as of today, the industry remains focused on the immediate "Resilience" mandate. Under the new Statement of Strategic Priorities (SSP), Ofcom is now legally tasked with ensuring that the business connectivity market remains competitive and secure.

For UK networking professionals, February 2026 is the month the "dumb pipe" finally died. Whether through the massive consolidation of metropolitan fibre or the arrival of autonomous 5G SA cores, the network is becoming a self-aware, software-defined ecosystem. As we bridge the gap between 5G maturity and the first 6G pilots, the priority is clear: the underlay must be as intelligent as the applications it carries. ■

IN DEPTH: Co-locate p10

Human error is costing UK businesses billions as AI cyber-attacks continue to soar

Whilst many businesses invest in state-of-the-art computers and up-to-the-minute software to keep them safe online, the digital experts at West Lothian's Workflo Solutions are warning the most exploited vulnerability for many businesses is actually human error. And so to mark CyberScotland Week, which runs from Monday 23rd until Saturday 28th February, they are urging everyone to tighten up their protection.

Cyber-attacks that target small to medium-sized businesses have increased sharply in recent years, propelled by advances in AI technology. From phishing scams and ransomware to credential theft and web-based intrusions, hackers are harnessing the power of AI to increase the speed, scale and sophistication of these devastating online attacks, stealing sensitive data for financial gain with greater ease than ever before. And according to Workflo Solutions, many employees are simply not keeping pace with these fast-moving developments. As a result, they are putting businesses at significant risk.

It is estimated that more than 95% of data breaches in the UK are caused by human error. Significant incidents – categorised as those costing at least £500 – average nearly £195,000 per business, amounting to a staggering £14.7billion worth of damages per year to the UK economy.

The most common human errors include falling for phishing scams, using weak or compromised credentials, misconfiguring security settings, and falling victim to social engineering.

Commenting ahead of CyberScotland Week, Workflo Solutions' managing director Michael Field urges businesses to stay vigilant; "CyberScotland Week is an opportune time to remind businesses of the need to create a strong security posture through their staff:

Create a culture where cyber security feels simple, relevant and part of everyday work.

Keep training short, frequent and practical by using real examples of phishing, social engineering and reconnaissance that staff might encounter.

Reinforce good habits through reminders, clear policies and easy routes to report suspicious activity.

Lead by example and reward secure behaviour so employees feel responsible and motivated to protect the organisation.

As AI becomes more integrated into everyday systems and business processes, the overall attack surface continues to grow, creating new entry points for exploitation within interconnected digital environments. It is therefore vital that we come together during CyberScotland Week to raise awareness, boost resilience and reduce the risk of human error."

To give businesses an extra helping hand during CyberScotland Week, Workflo Solutions is offering 200 businesses a free dark web domain audit. By heading to www.workflo-solutions.co.uk and signing up for this free service, businesses will be notified if any of their company email addresses or passwords have been breached, enabling them to take proactive steps to reverse these vulnerabilities and avert a potential damaging attack. ■

Suffolk charities benefit from £50,000 from CityFibre to boost digital inclusion

Suffolk Community Foundation has been awarded £50,000 from CityFibre's Community Fund to invest in its Digital Inclusion Fund. This significant grant addresses digital barriers across Suffolk and reflects CityFibre's ongoing commitment to empowering communities for a full-fibre future alongside Project Gigabit contracts. These community initiatives will help residents unlock opportunities through a shared goal: enhancing digital literacy, accessibility, and critical skills to thrive in an increasingly connected world.

Since its inception, the Digital Inclusion Fund has benefited multiple local organisations. Some key examples include the following groups.

- Headway Suffolk, a registered charity that provides support to those readjusting to daily life after a brain injury. Their share of the grant will be used to support their Digital Neuro Project, which delivers bespoke digital skills courses specifically designed for individuals living with a brain injury.
- Access Community Trust (ACT), which takes a holistic approach to end homelessness and tackles isolation by helping vulnerable individuals reintegrate online through their Connect and Thrive Programme.
- Eyes Open CIC, providing young people with subsidised training at their dedicated digital hub, The Vault. Having enabled over 900 placements with 28 schools, this funding ensures further tech-based sessions to combat youth isolation.
- River Church, which equips vulnerable Ipswich residents with vital Wi-Fi, laptops, printing access, and training to safely navigate the web independently.

Hannah Bloom, CEO at Suffolk Community Foundation noted that CityFibre's contribution has been absolutely instrumental. Since April, the Fund has supported ten organisations with grants to reduce digital poverty, improve skills, and increase device access. She highlighted a fantastic outcome: supporting more than 1,500 people across the county.

Additionally, CityFibre awarded a Small Grant to Ipswich. Love CIC. This funding helps maintain and expand their popular community platform, further fostering digital inclusion, civic pride, and local engagement.

Stacey King, Senior Partnership Manager at CityFibre, emphasized the importance of grassroots support to remove barriers. By investing here, they contribute to a more inclusive digital society where everyone can connect, learn, and thrive with confidence. ■

IPTel Solutions selects highlight to strengthen managed services strategy

Highlight announces that IPTel Solutions, one of Australia's leading wireless and enterprise networking businesses, has selected the Highlight Service Observability platform as a core component of its managed services strategy.

IPTel is one of Cisco's leading providers in Australia, delivering Cisco and Meraki Wi-Fi solutions to customers across the region. The integration of Highlight marks a significant step in IPTel's ongoing investment in its managed services capability and future growth plans.

Highlight provides IPTel with the ability to present customers with a clear, business-relevant managed service offering, underpinned by comprehensive monitoring, alerting and service visibility. Customers can choose to have IPTel fully manage the service, or they can also use the Highlight platform internally to gain greater insight into their own environments.

The platform was deployed rapidly, enabling IPTel to meet ambitious timelines. The extensive deployment took less than two months from the initial evaluation to go live. Highlight is now operational across IPTel's own infrastructure and its customer environments, leveraging Highlight's strong multi-tenancy capabilities to oversee multiple organisations from a single platform.

Costa Athanassiou, Chief Operating Officer at IPTel Solutions says, "We wanted to increase visibility for both us and our customers as we expanded our managed services practice. We had been using LogicMonitor for some time, but as that contract came to an end, we were looking for an alternative. LogicMonitor worked well for our engineers, but isn't as user friendly as we'd like.

"What sets Highlight apart is its ability to present technical information in a way that resonates with senior decision makers. The data is structured around locations and services, not just physical infrastructure with an intuitive look and feel which is ideal for CIOs, CTOs and C-suite stakeholders. These people are not interested in raw monitoring data; they want clear visuals that show how a branch or office location is operating."

IPTel's managed service philosophy is built on offering transparency to its customers. "We don't want to hide behind the curtain," adds Costa. "We want to be open and transparent about the services we're providing. Highlight allows us to show customers what's happening, have more meaningful conversations, and demonstrate the broader value of our managed services in a way that builds trust." ■

90 sites, one network: Kingspan's global connectivity revolution

Expereo, a global managed Network as a Service (NaaS) provider, announces that Kingspan Insulation—the manufacturer of premium insulation products—has selected it to deliver a fully managed connectivity solution across its worldwide operations.

Expereo's solution enables Kingspan to consolidate its complex network of disparate ISPs into a single, centralised platform. This enhances visibility, compliance, and agility to enable rapid growth and scalability.

Managing a complex WAN environment across 90 global sites with multiple vendors created significant challenges in efficient connectivity, compliance, and security. As the business moves away from MPLS towards SaaS and cloud services, the need for resilient, standardized, and agile connectivity has become critical.

Choosing Expereo as its global partner gives Kingspan the agility to scale quickly during acquisitions or streamline services when needed. Kingspan benefits from the design, implementation, and management of its solution, gaining a global view through expereoOne, Expereo's NaaS platform.

Through expereoOne, Kingspan now has full visibility and control over its network: from ordering and billing to ticketing and auditing. The portal also offers real-time performance insights, helping Kingspan optimise resources and maintain security.

Expereo now manages Kingspan's underlay and overlay services, including SD-WAN, Dedicated Internet Access (DIA), and broadband, with Low Earth Orbit (LEO) satellite connectivity under consideration for rapid manufacturing deployments. Remote SD-WAN setup completes in under 30 minutes, providing flexibility for urgent business needs without compromising performance.

This unified strategy underpins Kingspan's commitment to a single global standard. With a lean technical team managing 90 facilities across 28 countries, standardization is essential for driving efficiency. Expereo streamlines onboarding new locations and ensures seamless connectivity with automated failover.

This partnership follows a similar recent announcement from Kingspan Light, Air + Water, which also selected Expereo to deliver a fully managed connectivity solution for its worldwide operations. ■

PUBLISHING:
Publishing Director: Kathy Moynihan
Designer: Ian Curtis
Sub-editor: Gerry Moynihan
Contributors: Irvin Shillingford, Martin Smillie, Haris Pylarinos, Matt Smith, Kevin Cochrane

ADVERTISING & PRODUCTION:
Associate production:
Denyan Media Limited
Administration:
Karen Bailey

Networking+ is published monthly by: Kadium Ltd, Oatlands Mere, Weybridge, Surrey, KT13 9PD, Tel: +44 (0) 1932 886 537

ISSN: 2052–7373

The £1.6bn bold bet: AI's role in public service reform

UK public research funder, UKRI, has set out its first-ever plan to put AI to work for the nation's world-leading innovators. Building on historic strengths in computing and agentic AI, the strategy aims to transform cutting-edge research into innovations that benefit us all—from better healthcare to new public services and cutting-edge products.

UKRI-backed work is already making a difference through the RADAR AI system, which detects railway faults in real time, and the IXI Brain Atlas, supporting over 40 clinical trials for Alzheimer's. In the recent Spending Review, UKRI committed a record £1.6 billion directly targeted at the AI sector over the next four years, its biggest single investment area for 2026 to 2030.

Deputy Prime Minister David Lammy, leading the UK delegation at the India AI Impact Summit, noted: "From spotting cancers earlier to cutting backlogs in public services, new research into AI will be a game-changer." UK AI Minister Kanishka Narayan added that the plan will "harness AI to accelerate both the pace and possibility of scientific endeavour."

The strategy commits to expanding doctoral and fellowship routes co-designed with businesses, supporting high-paying jobs for research software engineers and data scientists. Under this new strategic framework, investment will focus on six priority areas: advancing technology development; transforming research through AI; developing AI skills and talent; accelerating innovation for economic growth and societal benefit; championing responsible and trustworthy AI; and building world-class AI data and infrastructure.

Professor Charlotte Deane, Executive Chair of the EPSRC, emphasized the need for bold choices: "UKRI will play a central role in backing the full innovation pathway from fundamental research to prototypes to scale-up." Immediate deliverables include £137 million for AI-enabled drug discovery and new treatments, plus £36 million to upgrade the University of Cambridge's "DAWN" supercomputer.

Today's AI is built on decades of curiosity-driven research. This strategy ensures discovery-led research thrives by simplifying programmes and removing barriers to scale. However, while this roadmap provides the financial engine, its success depends on the regulatory environment. As the focus shifts to deployment, the industry now looks to the government's newly finalized Statement of Strategic Priorities (SSP). This statutory framework will determine how UKRI-funded research is integrated into the nation's physical and digital infrastructure. ■

Agentic ops: 91% of security teams pivot to autonomous defence

Ivanti has announced the findings of its 2026 State of Cybersecurity Report: bridging the divide. Drawing on insights from over 1,200 global professionals, the report highlights a rapidly widening gap between escalating cyberthreats and organisational defence capabilities. AI is actively reshaping this landscape. Defenders remain optimistic, believing they use AI more effectively than threat actors—a confidence expected to grow significantly over the next two years.

However, operationalising this technology remains limited. While 93% of UK respondents state automation reduces response times, AI adoption across critical functions varies: 56% utilise AI for cloud security, and just 45% for threat intelligence. Trust in autonomous agentic AI is growing cautiously, with 91% making its integration a priority.

Daniel Spicer, Chief Security Officer at Ivanti, warns of a Cybersecurity Readiness Deficit: a persistent, widening imbalance in an organisation's ability to defend its networks. He notes that while AI effectively handles routine cyber hygiene, businesses are falling behind the evolving threat landscape due to rapid technological changes and SaaS initiatives.

The research identifies several key vulnerabilities. There is a growing IT-security rift: 39% of professionals feel IT lacks urgency regarding cyber risks. Furthermore, organisations take a fragmented approach to metrics: relying heavily on process-driven data rather than business impact analysis.

Deepfakes also present a surging threat: an alarming 76% of UK organisations have faced these attacks, with half experiencing sophisticated phishing emails. Yet, only 32% believe their CEO could reliably identify a deepfake. Finally, workforce burnout creates a silent systemic vulnerability. With 76% reporting that work stress harms their health, a lack of skilled talent remains the primary barrier to excellence.

This stark contrast underscores the urgent need for better strategies: balancing risk awareness with effective preparation to protect crucial networks. ■

Hiring 36 months ahead: The end of reactive recruitment

With critical technology skills in ever shorter supply, reactive hiring is no longer enough to sustain digital ambition. This feature examines how forward-thinking organisations are rebuilding talent pipelines — blending strategy, flexibility and automation to secure the expertise they will need long before demand hits.

In the middle of a chronic IT skills shortage, most organisations are still hiring like it's 2015.

Roles are approved when projects stall. Job titles are dusted off from outdated frameworks. Recruiters are briefed at the point of panic — not strategy.

The result is a familiar cycle: inflated salaries, rushed appointments, fragile teams and a slow accumulation of technical debt — not just in architecture, but in capability.

In today's market, the challenge is no longer finding people quickly. It is building a talent pipeline that scales before things break.

For Michael Aspinall, Senior Account Director at First Point Group, the defining feature of scalable pipelines is timing. "Effective pipelines are built well in advance, not in response to delivery pressure," he says. "They are grounded in a clear understanding of real skills and capabilities rather than inflated or outdated job titles."

That distinction matters. As platforms evolve and roles blur, job titles have become increasingly unreliable proxies for capability. Two "cloud engineers" may share a label — and almost nothing else.

The strongest pipelines, Aspinall argues, are deliberately blended. Permanent hires provide continuity. Contractors inject speed and niche expertise. Project-based talent absorbs peaks without distorting long-term structures. Equally important is where organisations look.

"Geographically diverse pipelines and talent drawn from adjacent industries are now essential," he says. "Ongoing talent mapping and market intelligence allow organisations to hire ahead of demand rather than chase it." In other words: stop fishing in the same shrinking pond — and start planning before the water runs out.

Not all shortages are created equal. The most difficult roles to pipeline at scale are no longer entry-level developers or generalists, but senior specialists with real-world operational scars.

"Platform and cloud engineers with exposure to large-scale, mission-critical infrastructure remain in particularly short supply," Aspinall notes. "In cybersecurity, the biggest gaps sit where technical risk must be translated into commercial decision-making."

Data and AI have shifted too. Demand is moving away from pure data science towards engineering, deployment and governance — narrowing talent pools that were already thin.

Meanwhile, emerging roles in AI infrastructure and MLOps are evolving faster than universities, training programmes and employers can respond.

The uncomfortable implication? Short-term hiring plans are obsolete.

"In practice, organisations need to be planning 18 to 36 months ahead," says Aspinall, "combining long-term pipelining with access to specialist contract capability."

For leaders still planning quarter by quarter, that timeline lands somewhere between inconvenient and existential. ■

Revolutionising the customer experience for the Welsh gigabit race

Route 101, a prominent systems integrator and cloud contact-centre specialist, has successfully implemented an integrated customer experience platform for Ogi, Wales's principal broadband provider.

This project has unified Ogi's contact centre operations, improved reporting capabilities, and streamlined digital engagement channels, enabling faster and more consistent customer service nationwide.

Ogi delivers ultrafast broadband to homes and businesses across Wales, with its contact centre handling inquiries related to installation, technical support, billing, and telesales. Previously, legacy systems hindered the company's ability to offer a comprehensive, omnichannel customer view, relying instead on separate tools for voice, social media, and email communications.

To address these challenges, Route 101 deployed the NiCE CXone Mpower platform integrated with Gamma SIP Connectivity and connected it seamlessly with Microsoft Dynamics CRM.

The new, fully integrated solution consolidates voice and digital channels into a single interface, empowering agents with access to complete customer histories, enabling quicker issue resolution, and ensuring consistent engagement across all communication touchpoints.

Alexander Breverton, Director of Sales Operations at Ogi, praised the implementation team's dedication: "The Route 101 team was highly engaged and attentive, often going beyond the original scope. From pre-sales to delivery, they provided excellent service and attention to detail."

Key outcomes of the project include:

A unified contact centre platform integrating telephony, social media, and reporting tools Streamlined management of social channels such as Facebook, Twitter, and LinkedIn Reduction in customer wait times through smarter call routing and improved first-call resolution Enhanced analytics and reporting to support ongoing service improvements Russell Attwood, CEO and Founder of Route 101, expressed enthusiasm about the project: "We're proud to support Ogi in building a future-ready customer experience platform. By unifying channels and embedding smarter analytics, we've helped them deliver exceptional service and position for future digital growth."

Looking forward, Ogi plans to expand platform capabilities with features like auto-summary to reduce agent after-call work and interaction analytics to identify coaching opportunities and further optimize service delivery. ■

The green AI superpower: efficiency as a competitive edge

The UK AI trade association, UKAI, is urging the nation to become the premier hub for Green AI—an approach prioritizing efficiency, cost-effectiveness, scalability, and sustainability.

Future AI leadership relies on effectively integrating the technology into energy systems and infrastructure. The race for dominance is shifting from raw computational power to delivering maximum value per energy unit, prioritising broader economic and social benefits.

UKAI notes the UK is uniquely positioned to pioneer this shift. High energy costs and grid constraints serve as catalysts for innovation in efficiency and system design. To secure leadership, UKAI advocates four priorities: integrated infrastructure aligning energy and digital networks, fairer pricing models for data, targeted innovation enhancing AI resilience, and smarter adaptive systems to optimize resources.

UKAI CEO Tim Flagg stressed the urgency: "the UK is at a crossroads; we have the chance to become the next green AI superpower if we act quickly." He calls for immediate action to leverage the UK's robust research base and operational experience.

Cybersecurity remains vital. Charlotte Wilson, Head of Enterprise UKI at Check Point Software, emphasizes that energy-efficient AI must have security at its core. With threats against utilities rising, AI deployments require robust security by design to protect national infrastructure and ensure long-term growth.

Acting now positions the UK as a trailblazer in sustainable AI, delivering economic, security, and geopolitical advantages for years to come. ■

Word on the web...

Rebuilding cloud control for the next era of UK enterprise compute

Kevin Cochrane, CMO at Vultr

To read this and other opinions from industry luminaries,

visit www.networkingplus.co.uk

Scale smarter: How VeloxServ colocation powers secure, cost-effective infrastructure growth

As digital transformation accelerates across every sector, businesses are generating and processing more data than ever before. From SaaS platforms and e-commerce operations to financial services and media delivery, the need for secure, scalable, high-performance infrastructure has become business critical. For organisations that want control over their hardware without the capital burden of building their own facilities, colocation remains one of the most strategic infrastructure decisions they can make.

VeloxServ's colocation services are designed specifically to support ambitious, scaling businesses. By combining enterprise-grade datacentre environments with commercial flexibility and hands-on expertise, VeloxServ enables organisations to grow their infrastructure in line with demand – without compromising performance, security, or budget.

Building for Growth Without Building a Datacentre

Constructing and maintaining a private datacentre is costly, complex, and increasingly unnecessary. Real estate, power provisioning, cooling systems, redundancy design, security compliance, and ongoing maintenance represent significant capital and operational expenditure. For many organisations, those resources are far better invested in product innovation and customer growth.

VeloxServ provides a smarter alternative.

By colocating infrastructure within VeloxServ's secure facilities, businesses gain access to resilient power, cooling, connectivity, and physical security – without the overhead of building and managing it themselves. This enables organisations to retain full control of their servers, networking hardware, and storage, while benefiting from a purpose-built environment engineered for uptime and performance.

Enterprise-Grade Security and Resilience

In today's threat landscape, infrastructure security is non-negotiable. VeloxServ's datacentre environments are ISO27001 certified, demonstrating a rigorous approach to information security management. For customers, this provides confidence that physical access controls, environmental safeguards, and operational processes meet internationally recognised standards.

Resilience is equally critical. Redundant power feeds, robust cooling systems, and carrier-diverse connectivity ensure that customer infrastructure remains available and performant. Whether supporting latency-sensitive applications, transactional systems, or high-throughput platforms, VeloxServ's facilities are engineered to keep businesses online and operational.

For organisations operating in regulated industries or handling sensitive data, this combination of compliance and reliability provides a strong foundation for growth.

Scalable Solutions for Evolving Demands

One of the primary advantages of colocation is scalability. Infrastructure requirements rarely remain static. A product launch, geographic expansion, customer acquisition surge, or data analytics initiative can all drive rapid increases in compute and storage needs.

VeloxServ's colocation solutions are built to scale alongside customers. Businesses can begin with the footprint they need today – whether that is part rack, single rack, or multiple racks – and expand seamlessly as demand increases. Flexible power densities and adaptable configurations ensure that growth does not require disruptive relocations or costly redesigns.

For growing organisations, this scalability provides strategic agility. Infrastructure can evolve in step with commercial objectives rather than becoming a limiting factor.

Expert Support from Planning to Deployment

Colocation is not simply about rack space; it is about partnership.

VeloxServ supports customers throughout the entire lifecycle of their deployment. From initial capacity planning and power specification to rack layout optimisation and connectivity design, the VeloxServ team works closely with businesses to ensure their infrastructure is implemented efficiently and correctly.

For companies migrating from on-premise server rooms or legacy facilities, this guidance is invaluable. Transitioning infrastructure requires careful planning to minimise downtime and avoid configuration errors. VeloxServ's experience reduces that risk, enabling smoother deployments and faster time to value.

Craig Messer, Managing Director of VeloxServ, explains: "Businesses today can't afford for infrastructure to slow them down. Our goal is to remove complexity and give customers the flexibility to scale quickly, securely and cost-effectively. Whether they're deploying a single rack or expanding a large estate, we work as an extension of their team to ensure everything is delivered right first time."

Ongoing technical support ensures that as customer environments evolve, assistance is always available to optimise performance and accommodate change.

Reducing Upfront Costs: A Limited-Time Colocation Incentive

Recognising that flexibility and cost efficiency are more important than ever, VeloxServ is launching a limited-time colocation incentive designed to accelerate customer growth while lowering initial investment.

From February 2nd, all customers taking new or additional colocation space will benefit from three months of free billing. Whether expanding an existing footprint or deploying into VeloxServ's datacentres for the first time, organisations can realise immediate cost savings while still accessing enterprise-grade performance and security.

This incentive directly addresses one of the most common challenges businesses face when scaling infrastructure: upfront expenditure. By removing three months of billing, VeloxServ enables companies to reallocate capital into hardware investment, product development, recruitment, or market expansion.

For customers committing to a full rack or more, VeloxServ is extending the offer even further. The company will support the transport of equipment into its datacentres at no additional cost. Logistics can often be a hidden complexity in infrastructure projects, particularly when moving sensitive hardware. By absorbing this burden, VeloxServ simplifies deployment and reduces friction at a critical stage of scaling.

Enabling Confident Infrastructure Expansion

In a competitive digital landscape, infrastructure cannot be an afterthought. Performance impacts customer experience. Security protects brand reputation. Scalability underpins innovation.

VeloxServ's colocation services empower businesses to expand confidently, knowing their infrastructure is housed within a secure, resilient, and professionally managed environment. The current limited-time incentive further strengthens the business case, offering immediate cost relief and logistical support precisely when it matters most.

For companies considering expansion or evaluating a new colocation deployment, there has rarely been a better time to act.

To discuss your requirements and secure your free billing period, contact the VeloxServ team today at sales@veloxserv.co.uk ■

Ransomware attacks: building a human firewall to defend against the rising threat

Irvin Shillingford, regional manager Northern Europe, Hornetsecurity.

2025 saw a surge in cyberattacks across nearly every industry ranging from car manufacturing to luxury fashion, few sectors have escaped unscathed. Last year marked a critical turning point, with our annual Ransomware Impact Report identifying the first rise in ransomware attacks in three years. This signals the end of a period of decline and the beginning of a renewed, intensified wave of cybercrime.

The report reveals that nearly a quarter (24%) of businesses fell victim to ransomware attacks last year, a sharp increase from 18.6% in 2024. As businesses navigate this cyber landscape, the findings highlight the need for continuous employee education, investment in advanced security measures, and the integration of AI to combat sophisticated cyber threats. The rise demonstrates why it's essential for business leaders to fully understand how they can avoid falling victim to a ransomware attack.

The risk has never been higher

The resurgence in ransomware incidents is a warning for organisations across all sectors. The increasing number of businesses being targeted underscores how cybercriminals are becoming more advanced, organised and persistent. Last year alone, NatWest bank described the "continuous arms race" as its systems are targeted by around 100 million cyberattacks every month. This threat landscape is characterised by global groups who are targeting specific vulnerabilities of technology and human behaviour and acting systematically to exploit them.

While the rise in ransomware attacks is a warning cry to businesses, it seems they are not taking sufficient proactive steps. Only 46% of businesses are insured against ransomware attacks last year, compared to 54.6% in 2024. This decline is linked to more stringent requirements on the part of insurers and yet it comes at a time when the financial stakes are escalating dramatically.

Global ransomware payments reached a significant estimate of $1.1 billion in 2025, a figure amplified by the use of AI-driven attack methods. Advanced spoofing techniques now use legitimate brand email addresses in ways that can elude detection, and AI can mimic the language of colleagues to make the communications seem even more legitimate.

Analysts recently named the Jaguar Land Rover hack the costliest cyber-attack in UK history, with the impact of their manufacturing shutdown having a domino effect throughout their supply chain and the economy. Robust insurance, coupled with a strong security strategy and proactive tactics, is essential not only for businesses to protect themselves but also to keep afloat.

Your employees are assets with potential waiting to be unlocked

As cyberattacks become more frequent and complex, traditional forms of defence are no longer sufficient. It is crucial for organisations to protect their people and their data with multi-layered, robust, AI-powered solutions. As part of this, immutable backup storage is a must as part of any sound Disaster Recovery Plan.

Tech is a critical defence, but people are too. One wrong click can bring a business to a halt, making the human firewall of employees trained to recognise and respond to cyber threats just as vital as any technology.

Security awareness training continues to be a key pillar of ransomware defence. Almost three-quarters of organisations now provide some form of end-user training to counter threats, recognising that cybersecurity is not the sole responsibility of the IT department, but a shared obligation across all levels of an organisation. Employees of every skill set and seniority must undergo comprehensive, ongoing education to remain alert to evolving tactics.

Building this human firewall requires sustained commitment. Training programmes must be ongoing, relevant, and tailored to each individual to keep pace with the changing threat environment. This approach transforms employees from potential weak points into proactive defenders who strengthen an organisation's overall resilience.

AI and advanced training are essential

Organisations must adopt next-gen solutions that combine the strengths of AI and adaptive learning technology to create a strong security culture.

Although 74% of organisations report offering end-user training against ransomware, 42% of security leaders admit that their training is insufficient. AI can address this gap by enabling personalised learning experiences that are both scalable and effective. Through simulation-based exercises such as regular spear phishing tests, employees gain hands-on experience in recognising and responding to realistic attack scenarios. Such training helps to prevent "false compliance" - a superficial level of cybersecurity awareness with low-quality, check-the-box training which lacks adequate follow-up and contributes to ongoing human error.

AI-powered training platforms offer a transformative approach to cybersecurity education. These systems ensure that employees receive the most relevant and timely training possible by automating the delivery of content and tailoring it to each individual's behaviour and role.

Businesses need to protect themselves

There can be no room for complacency. Businesses should have strong cybersecurity tech provisions to address ransomware across the board by having the right security infrastructure, appropriate insurance, building a strong human firewall to mitigate the risk of costly breaches, and leveraging AI-driven solutions that personalise training.

Integrating AI with a culture of cybersecurity awareness will enable businesses to transform their workforce into a security asset against the growing wave of ransomware attacks. ■

Your infrastructure changed 50 times today. Your compliance check didn't.

Consider Klarna's reality: a fintech organisation with 30+ engineering teams, each managing its own cloud infrastructure. Every team deploys multiple times daily: new services, updated configurations, additional regions.

Consider Klarna's reality: a fintech organisation with 30+ engineering teams, each managing its own cloud infrastructure. Every team deploys multiple times daily: new services, updated configurations, additional regions. The cloud inventory team attempted to maintain a master view of security controls across all these deployments using the traditional approach. The math simply broke. By the time they finished validating one quarter's infrastructure, hundreds of changes had already happened in the next. The compliance snapshot was outdated the moment it was complete.

Klarna's response transformed how they think about infrastructure governance. They created over 100 "target specifications" covering nearly every aspect of configuration management — from database operations and tagging to logs and encryption. These target specifications define expected states as code and continuously validate against actual infrastructure. The results: violations are detected immediately rather than periodically, and teams can self-correct before issues compound. The example at Klarna provides a window into infrastructure at scale.

The velocity gap

Platform engineering teams deploy infrastructure changes 50+ times per day as standard practice. Infrastructure-as-code has transformed network configurations into software, creating environments filled with ephemeral resources spanning multiple clouds and automated pipelines. Each deployment can modify firewall rules, network segmentation, IAM policies, or encryption configurations.

A central European bank might have 100+ engineering teams operating across five cloud regions. Periodic validation cycles cannot keep up with thousands of daily configuration changes. The gap between "approved state" and "actual state" grows wider by the day, posing a dangerous challenge. As your organisation passes compliance audits, the actual security posture steadily degrades.

DORA requirements demand proof of operational resilience, not just the existence of a policy. NIS2 has expanded its scope to cover critical infrastructure, supply chain visibility, and incident reporting requirements. The penalty structure reaches up to 2% of global revenue for non-compliance — high stakes that demand attention.

What regulators demand represents a fundamental shift in how organisations must demonstrate compliance. The requirement has evolved from "we have a firewall policy" to "we can prove every firewall configuration matches policy right now." Organisations must prove those controls work continuously. Networking teams are on the hook because they implement these controls in practice.

First audit cycles under these new frameworks will reveal an uncomfortable truth: their validation approaches are fundamentally insufficient.

Traditional validation is broken — the truth of what continuous really means

Point-in-time audits are snapshots that age like milk. Manual evidence collection — network diagrams, screenshots, interview notes, spreadsheet exports — documents what existed at a specific moment, not what exists now.

Hybrid complexity compounds the problem exponentially. Organisations run on-premises infrastructure alongside AWS, Azure, GCP, and numerous third-party services. Consider a common requirement: MFA must be enabled and enforced for all privileged access. Multiple control tests are needed to validate this properly — verifying MFA is enabled across identity providers, checking enforcement policies for administrative roles, and confirming exception processes are documented and approved. Each of these elements can drift independently, across different systems, modified by various teams. Security teams physically cannot manually review every infrastructure change. When validation cycles can't keep pace with deployment velocity, configuration drift becomes inevitable.

Continuous validation isn't about conducting more frequent audits. It's about embedding validation directly into operations, treating infrastructure testing like software testing rather than as a separate compliance exercise. Rather than validating controls after deployment, organisations now catch misconfigurations at commit time — before infrastructure even reaches production. Automated monitoring continuously detects drift in real-time: a firewall rule no longer matches the approved baseline, 15 S3 buckets suddenly lack encryption, or a new service violates zero-trust boundaries. This shift from reactive auditing to proactive validation means issues surface immediately across multi-cloud environments, network configurations, and encryption standards, allowing policy violations to be caught before they become operational problems.

The "shift-left" principle catches control violations before they reach production environments. Real-time evidence collection means every validation automatically generates an audit trail. By integrating validation into CI/CD pipelines, validation becomes part of the deployment process rather than a separate activity. The key difference: validation happens continuously in the background, not as a scheduled event that disrupts operations.

From quarterly audits to continuous confidence

Start with the highest-risk controls: encryption standards, network segmentation, privileged access. Define control tests as code alongside infrastructure definitions, integrating them into existing CI/CD pipelines rather than creating parallel processes. Build automated evidence collection directly into change workflows.

The cultural challenge is significant. Engineering teams must accept validation as inherent to deployment, not an obstacle. Start small — pick one critical control, automate its validation, and prove the value. Then scale gradually, adding controls as teams adapt to the continuous validation model and build confidence in the approach.

Security validation must become part of infrastructure operations, not a separate compliance exercise. This isn't about slowing down — it's about moving fast with confidence.

The tools exist. The barrier is organisational will. DORA and NIS2 are forcing this conversation, but the real driver is balancing business velocity with regulatory reality in increasingly complex environments.

Teams that figure out continuous validation turn compliance into a competitive advantage. Those still operating on quarterly cycles will struggle when regulators arrive, demanding continuous proof. ■

Building early warning alerts to reduce knock-on, escalating disruption

Martin Smillie, Senior Vice President, Communications and Data Exchange at SITA.

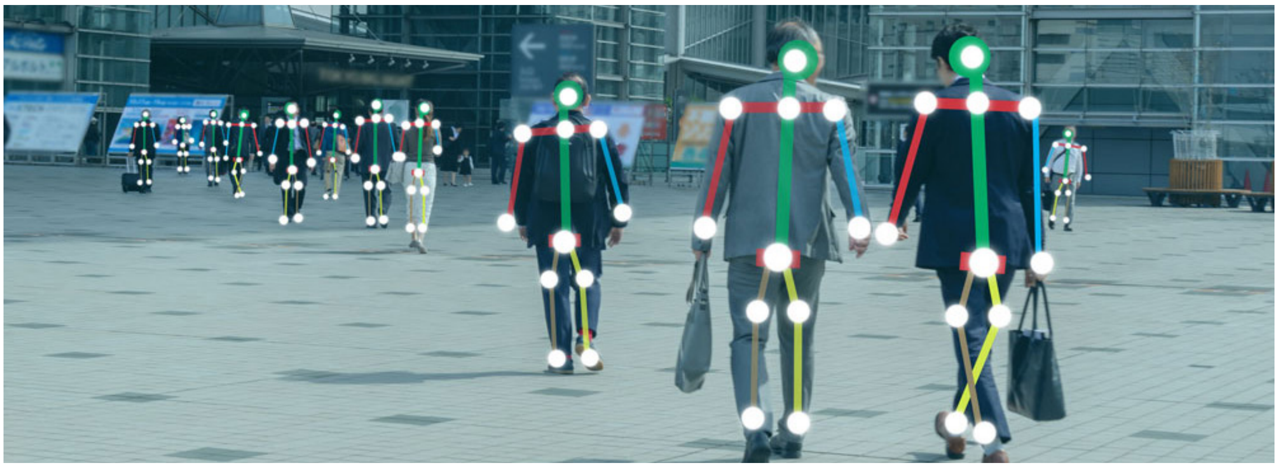

To quote Benjamin Franklin: "Every delay you do not see in time is a delay you no longer truly control." He realised the importance of being proactive and aware of our actions – because time is a precious resource. Indecision and delays simply divert attention from our original goals. It's where predictive analytics can now be used to positively influence an operational landscape.

In an industry built on precision, disruption in aviation and other modes of transport can still ambush you on the day of operations, turning a manageable late departure upstream into an avoidable mess at the intended point of arrival.

The real cost of late visibility

Most damage isn't caused by the primary delay itself, but by how long it takes to become visible to those expected to manage it. By the time an inbound delay finally lands in systems in a usable way, options may have quietly evaporated.

Ground teams have been rostered and are waiting at the wrong time. Gates are occupied by aircraft that will not push back as planned. Crews are edging toward the limits of legal duty.

Connection windows are shrinking and passengers are still walking to the next gate; unaware their departure schedule is dropping by the minute.

This is where the cost truly lives: in trapped resources, missed opportunities to reassign, and the reputational hit when an experience becomes one that feels chaotic and unfair.

Forced to react, not intervene

Operational control in that world is fundamentally defensive. Throw people at the problem, rework the schedule under pressure, and lean on heroics from frontline staff. But you rarely get what's really needed: time. Time to move a gate. Time to re-sequence turnarounds. Time to adjust crew plans before legal limits bite. Time to protect onward schedules and connections.

The irony is that information already exists somewhere in the ecosystem to give you time, but it's just not reaching teams, early and reliably enough, in a form they can act on.

Early warning, not early panic

What would it look like if transport treated delay intelligence the way other industries treat demand forecasting or supply chain risk – something you anticipate, model and act on before the impact hits?

The shift is subtle but profound: from "knowing a flight is late" to "knowing early that it is likely to be late, and by roughly how much." That distinction is where you win back control.

Predictive AI-generated analytics can then help us drive a proactive, actional posture that minimises any escalating impact on operations.

With these earlier warnings, ops centres are no longer scrambling to rebuild the plan. They're shaping it:

- Moving service teams and ground handlers to where and when they will be needed, not leaving them idle.
- Sequencing aircraft turnarounds with a realistic view of inbound flows, not a fragile, on-paper schedule.
- Making calm, deliberate decisions on gate allocation, rather than swapping stands under duress.
- Adjusting crew plans and connection handling with time to give passengers a fair chance, and their network a fighting chance.

Early signals are about creating the space for better judgment. And that can build greater profitability. IATA has estimated that air traffic management delays in Europe have cost airlines and passengers €16.1 billion over the past decade. So there's potential for a huge gain by getting better.

Flow is now your primary currency

Complexity has increased, but buffers have not. Schedules are tighter, aircraft and crews are more heavily utilised, and passengers are less tolerant of poor communication.

Flow is everything – of aircraft, of people, of information. Having invested in keeping physical flows moving through terminals, borders and turnarounds, the next competitive frontier is the flow of operational intelligence.

If data exchange is being bridged by shared platforms behind the scenes, the question is whether that data is being used early, and intelligently enough, to protect operations before disruption spreads.

From schedule worship to scenario thinking

For teams still treating the published schedule as a sacred truth: "right" until something proves it "wrong", the reality is the published schedule is more like a first draft.

Organisations that thrive are those that normalise continuous re-interpretation of that draft considering new signals. They build muscles around scenario thinking:

- If a departure slips by 15 minutes, what does that mean for the arrival?
- If a connection is at risk, which passengers, crews and rotations need protecting most?
- If a gate is going to be blocked, what's the lowest-friction way to re-sequence?

Earlier, reliable AI-driven insight simply pulls thinking forward in time.

A better deal for your passengers and your people

Passengers rarely distinguish between "your" delay and someone else's. They judge how clearly you communicate, how fairly you treat them, and how much of their time you waste.

Timely, honest expectations – "we regret this flight is now expected to depart 15 minutes later than schedule" – can make the difference between frustration and understanding. It also presents a chance to re-time passenger processing, adjust servicing and manage onward plans with empathy and precision.

Turning disruption into a leadership opportunity

This underlines why delay performance and resilience are increasingly seen as leadership metrics, not just operational ones.

We cannot eliminate disruption from aviation or other modes of transport, but we can minimise it. Weather, airspace congestion, and technical issues will surface. But we can decide if an airport or airline is systematically ready for disruption and utilising predictive analytics to address it.

Leaders who choose this path will not just protect their schedules; they'll redefine what "on-time" means in a connected, data-rich industry. They'll embrace earlier, smarter signals as a strategic asset, not an optional add-on. And they'll build an operation where delays may still occur – but rarely catch anyone off guard. ■

Cybersecurity risks rising for UK law firms in 2026

Matt Smith, Chief Technology Officer, Softwerx.

As 2026 gets underway, the cybersecurity landscape for small and medium-sized UK law firms is becoming increasingly challenging. Security incidents are no longer rare exceptions, they are a predictable challenge of running a modern legal practice. Firms that fall short in protecting client information and ensuring operational continuity face serious commercial and reputational consequences.

In this environment, resilience is paramount. The firms that can rapidly identify threats, limit disruption and maintain client confidence will be the ones best positioned to operate steadily when these challenges arise.

So, what does the coming year look like? This 2026 cybersecurity forecast explores the pressures already affecting UK law firms, the risks that are intensifying and the steps practices can take to remain secure, compliant and functional in a year where cyber incidents have become a routine business concern.

Today's forecast: heavy cloud is here to stay

For most UK law firms, day-to-day legal work now happens almost entirely in the cloud. Email, documents, collaboration platforms and more recently, AI-powered tools all sit within cloud-based productivity suites, making the cloud the central operating environment for modern legal practices. This shift has brought major benefits in flexibility, scalability and efficiency but it has also concentrated a significant amount of cyber risk in a single place.

As we move through 2026, a firm's security maturity is no longer measured by how many tools it has deployed. Today's threats are less about missing technology and more about misconfigurations, weak governance and a lack of continuous oversight. What matters now is how well those tools are set up, managed and monitored to maintain resilience and meet regulatory expectations.

Rising temperatures in the identity and email threat landscape

Phishing, impersonation and credential theft remain the most common entry points for attackers targeting law firms, but the nature of these threats is shifting rapidly. In 2026, cybercriminals are increasingly using AI to speed up their operations, automate large parts of their campaigns and craft messages that are far more convincing than those seen in previous years.

As a result, identity is where pressure builds first. When access rules are too permissive, multi-factor authentication is applied inconsistently or access privileges extend beyond genuine need, attackers no longer have to force entry or exploit complex technical weaknesses. They simply sign in through systems that assume trust, rather than verify it.

Warnings for the months ahead

A strong probability of incidents

Cyber incidents are no longer unusual events for UK law firms. They've become a routine part of the operating environment. In 2026, the critical question is not whether a firm will be targeted, but how quickly it can identify an issue, contain the damage and restore operations without long-term fallout. Even short periods of downtime can halt casework, interrupt billing cycles, threaten deadlines and cut off vital client communication.

Reduced visibility during AI adoption

AI is rapidly becoming embedded into everyday legal work. Tools for document summarisation, contract analysis and research are now widely used across many practices. AI does not introduce entirely new risks on its own, but it amplifies existing weaknesses in areas such as access control, data management and governance. Even when modern AI platforms are designed to support data-residency and privacy expectations, safe and compliant use still relies on strong identity, device and data-protection controls underneath.

Pressure systems across the sector

Insurer pressure is increasing

Professional indemnity and cyber insurance tailored for the legal sector are now closely tied to verifiable security practices. Insurers are no longer reassured by policy documents or stated intentions - they want concrete proof. In 2026, requirements such as multi-factor authentication across all access points, strong endpoint protection and patching, reliable and tested backups and a well-defined incident response plan will be baseline expectations. Renewal conversations are increasingly becoming comprehensive security evaluations in their own right.

Regulatory pressure is steady and unforgiving

Regulatory expectations will continue to tighten in 2026 with SRA standards, UK GDPR obligations and ICO enforcement all reinforcing the same message. Law firms are expected to protect client confidentiality, safeguard client funds, report incidents promptly and demonstrate reasonable preventative measures. When basic safeguards are missing, penalties and enforcement action follow.

Supply chain risk is persistent

Third-party access remains a quiet but persistent source of exposure and risk. Legal technology platforms, managed service providers and shared client environments all extend the attack surface. Firms that regularly review third-party access rights and actively monitor for unusual behaviour in the year ahead are far more likely to identify issues early, before they escalate into full incidents.

How to prepare before the climate turns

Most UK law firms already own the security tools they need. The priority for 2026 is not buying more technology but optimising existing cybersecurity investments.

Secure identity management should be treated as non-negotiable for law firms operating in the cloud. Multi-factor authentication needs to apply to every account and access route. Privileged access should be tightly controlled, with unnecessary rights removed wherever they are not essential. In a cloud-first environment, identity is the primary line of defence.

As AI adoption accelerates, reducing permission sprawl and eliminating shadow AI becomes urgent. Broad access that once seemed manageable and even acceptable becomes far riskier when AI tools can surface and reuse information at speed.

Sensitive and privileged data must be protected far more deliberately. Clear classification and labelling ensure protection follows the data itself, reducing the risk of accidental sharing outside the firm or beyond the appropriate team during everyday collaboration.

Devices used for client work remain a foundational control. Every laptop and mobile device should be managed, patched and protected. With hybrid working now the norm, endpoint security is no longer optional hygiene, it underpins secure access.

Email and collaboration platforms continue to be the main entry point for attacks. While staff awareness matters, it must be backed by strong, automated protection to stop phishing and impersonation attempts before they reach users.

Cyber incidents should be assumed, not treated as unlikely events. A clear response playbook should define who decides, who communicates and what happens in the first hours following a breach attempt.

Finally, continuous monitoring is now a baseline requirement. Threats do not operate only during office hours and most law firms don't have the resources to staff a skilled 24×7x365 security capability internally, particularly those with lean in-house IT teams. Outsourced Managed eXtended Detection and Response (XDR) solutions monitor security signals round the clock and respond to threats, providing better visibility, faster containment and reduced pressure on internal teams.

The long-range forecast into 2027

Looking ahead, AI adoption in law firms will continue to accelerate, driving efficiency but also magnifying existing security risks. At the same time, insurer scrutiny will tighten and regulatory enforcement is expected to remain consistent, leaving no room for lapses in governance or controls.

In one sense at least, the forecast is not gloomy; it is clear. Preparation and response capability now matter more than prevention alone. Law practices that optimise their security investment and reinforce it with 24×7 Managed XDR oversight will still face challenges but with the enterprise-class cybersecurity capability they deserve, they will detect threats earlier, contain incidents faster and maintain operations when less-prepared firms are forced to pause. ■

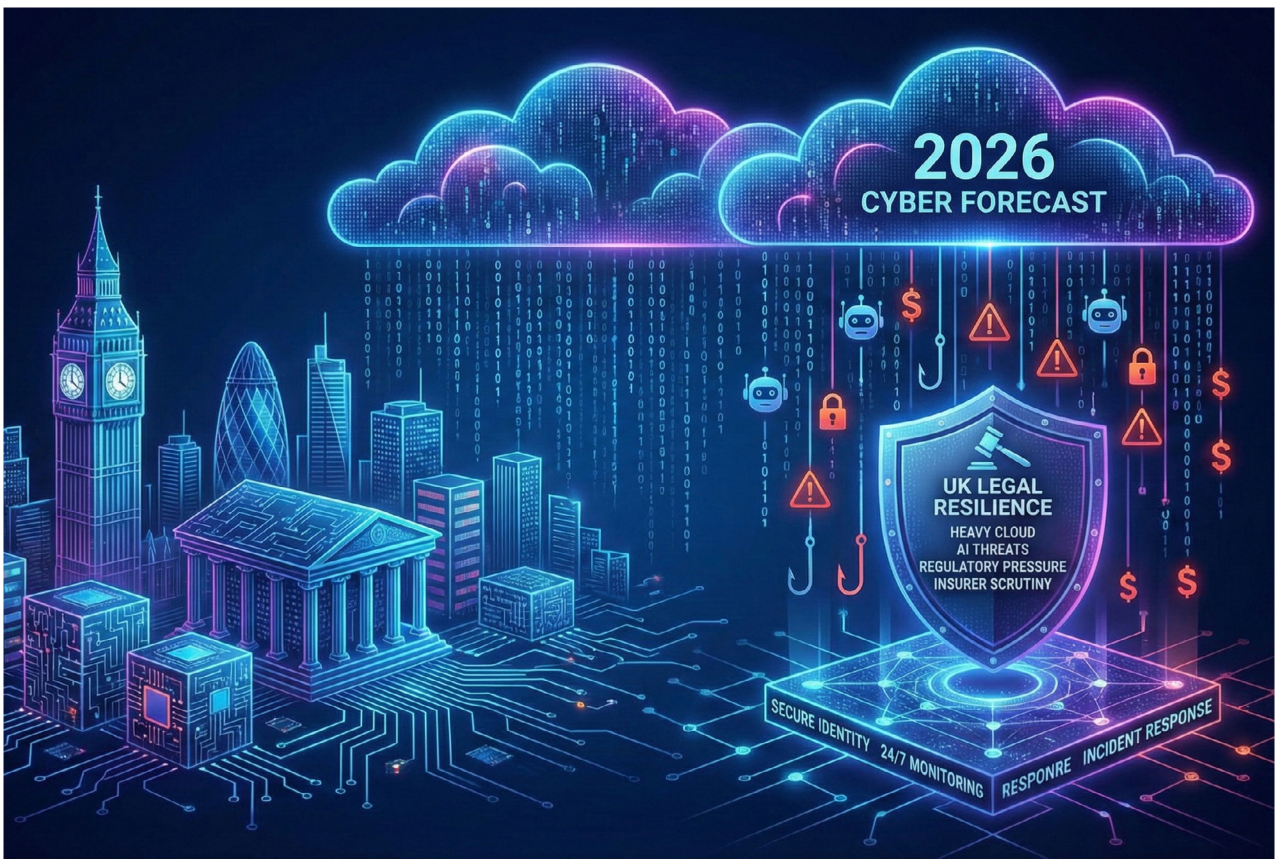

To co-locate or not to co-locate?

As UK enterprises rethink where their infrastructure should live, the old build-versus-buy debate is giving way to a far more complex set of choices around power, regulation, sustainability and flexibility.

For years, the question facing UK IT leaders seemed deceptively simple: build your own data centre or rent space in someone else's. In 2026, that binary choice has quietly collapsed.

Rising energy costs, sustainability pressure, power constraints and the normalisation of hybrid cloud have reshaped the conversation. Co-location is no longer just a cheaper alternative to bricks and mortar. Nor is on-premise infrastructure the stubborn relic it was once painted to be.

Instead, UK organisations are navigating a far more nuanced decision: where should each workload live, and why?

As James Sturrock, Director of Systems Engineering at Nutanix, puts it, "this is less about a binary choice and more about flexibility and outcomes."

Rethinking the build-versus-buy question

Traditionally, owning a data centre offered control, predictability and long-term financial logic — particularly for large enterprises with stable demand. But that equation is changing fast.

"Facilities have become far more sophisticated in terms of compute density, efficiency and scalability," says David Watkins, solutions director at VIRTUS Data Centres. "The question is no longer just cost, but performance, control and long-term strategic advantage."

Building a private data centre now requires not just capital, but deep technical commitment: specialist cooling, high-density power design, resilience engineering and continuous upgrades. For many organisations, that burden is increasingly hard to justify.

Stewart Laing, CEO of Asanti Data Centres, notes another critical factor: "Businesses need to think more about skills, OPEX and location, asking themselves where best they can access power and network, the latter of which can be extremely expensive depending on location to fibre backbone. For many, owning the hardware and putting it in a colocation data centre makes more sense, helping to shift location, power and connectivity risk to the data centre operator instead."

Co-location, by contrast, offers enterprise-grade infrastructure without the operational overhead. Providers invest heavily in R&D, certified staff and multi-site estates, allowing customers to scale, migrate and rebalance deployments across locations as needs evolve.

For Sturrock, the real differentiators are "cost predictability, operational complexity, resilience, and the ability to optimise IT resources while avoiding vendor lock-in."

It's not about where your servers sit, but about how intelligently you can move them.

Power, price and pressure: energy reshapes the decision

If there is one force tipping the scales towards co-location, it is energy.

With electricity prices volatile and grid capacity tightening in parts of the UK, running a private facility has become an increasingly specialist pursuit. Traditional architectures are power-hungry, and sustainability targets now sit squarely on the CIO's desk.

"Energy procurement is a multi-faceted, specialist process," says Watkins. "Data centre providers have dedicated teams securing long-term supply, often through Power Purchase Agreements directly with renewable generators."

These PPAs not only stabilise pricing, but actively increase renewable capacity on the market — something few single enterprises can realistically achieve alone.

Sturrock notes that modern hyperconverged platforms also play a role. "HCI-based architectures can significantly reduce energy demand and carbon footprint, which is why many organisations are re-evaluating where and how their infrastructure is run."

Laing highlights the broader challenge: "Tightening sustainability targets mean you can't stand still — you either fund upgrades every year or you lean on a colo that is already optimising power usage effectiveness (PUE) and renewables. What truly needs to be addressed for consumers and businesses alike is the cost of power. From a colocation provider perspective, this is as simple as being able to compete with our European counterparts. The uncomfortable reality is that if UK colocation facilities are outpriced by cheaper EU markets, it's bad news for our digital economy."

For organisations still operating legacy on-premise facilities, the challenge is stark: match hyperscale-grade efficiency, or accept growing operational and environmental penalties. (Spoiler: very few manage the first.)

Compliance without contortions

In regulated UK sectors — finance, healthcare, government — infrastructure decisions are inseparable from sovereignty and compliance.

Co-location can be a powerful ally. "It allows organisations to retain control, guarantee the physical location of data, and enforce residency policies," says Sturrock. "That's why we see strong demand from highly regulated sectors."

Laing concurs: "Colocation simplifies the model: the data centre provider owns the building and is responsible for the uptime of the infrastructure; the organisation owns the hardware, data, platforms and controls. Keeping workloads in UK data centres allows businesses to comply with sovereignty and residency regulations which can be more difficult in hyperscale public clouds, and ultimately gives them greater control."

Modern platforms now enable encryption, granular access controls and location-based policy enforcement regardless of physical site — offering compliance with far greater operational flexibility.

But co-location is not a magic wand.

"Complexity increases when governance isn't applied consistently across environments," Sturrock warns. Hybrid estates without unified policy often create more audit risk, not less.

And some workloads still resist migration. Highly sensitive data, ultra-tight security models and bespoke legacy systems often remain better suited to private environments.

Yet organisations persist in forcing them into co-location anyway — "due to legacy decisions, existing contracts, or risk-averse culture," rather than technical merit.

Laing adds nuance on workload suitability: "Very few workloads are fundamentally wrong for colo; the real constraint is whether you have the inhouse capability or partners to run your own kit properly. Early stage startups with no IT team often find it easier and cheaper to begin in public cloud but as part of their business planning, they should keep an eye on costs and eventually plan to move in-house by way of deploying into a colocation data centre."

Hybrid by default

Co-location is no longer a stepping stone on the way to "the cloud". Nor is on-premise staging a comeback. Instead, hybrid has become the default operating model.

"Hybrid multicloud is no longer a transition phase," says Sturrock. "It's becoming the standard operating model."

Watkins agrees. "Few organisations rely on a single infrastructure model today. Businesses are blending on-prem, colocation and hyperscale cloud to optimise resilience, cost and performance."

Laing frames the new role of colo: "Colo is now the anchor, not just a stepping stone. Organisations are increasingly buying their own hardware, installing it in colocation data centres and tapping into public cloud for burst capacity and R&D rather than as the default home for everything. It's less about cloud 'maturing' and more about customers becoming wise to cost and vendor lock in and using colo to regain control."

Increasingly, architectures span core, regional and edge facilities — with workloads placed where latency, regulation, cost and sustainability align best. Looking ahead five years, neither expert sees a single dominant model emerging.

Laing adds perspective on long-term trends: "Inhouse builds are likely to remain the preserve of a small number of operators running operation-critical, high value infrastructure like major banks and trading houses. A broad on-prem comeback looks unlikely given the cost and complexity of build, plus the need for continual reinvestment."

"What will define success is flexibility," says Watkins. "Interoperability will shape AI infrastructure strategies. The world has gone hybrid — and for AI, it's going multi-hybrid."

Which leaves CIOs with a deceptively simple mandate: stop asking where your data centre should be, and start asking where each workload truly belongs. After all, the future of infrastructure is not about buildings. It's about choices. ■

Trackside intelligence: What Anthropic's F1 entry means for the edge

The recent announcement that Anthropic will become the "official thinking partner" for Williams Racing, supported by Atlassian, marks a departure from the standard tech-sponsorship playbook.

While branding is part of the deal, the focus on engineers using Claude for trackside operations suggests a pivot in how AI companies want to be perceived.

For the data centre industry and the AI sector, this isn't just another sports marketing play. It is a real-world stress test for latency, edge computing, and the physical constraints of modern infrastructure.

From Generative Hype to Functional Utility

For the last two years, the AI narrative has been dominated by "magic"—the ability to generate text or images from thin air. However, the Anthropic/Williams partnership shifts the focus to operational utility. In an F1 environment, "good enough" data isn't useful. Decisions on tire degradation, fuel loads, and aero-elasticity happen in seconds.

By positioning Claude as a "thinking partner" for engineers, Anthropic is signaling that its models are ready for high-consequence, low-latency environments. This moves the needle for AI from being a back-office productivity tool to a core component of the industrial "stack." For data centre providers, this means the demand profile is changing. Clients are no longer just looking for bulk compute to train models; they need highly responsive, high-availability environments that can support real-time engineering feedback loops.

The Shift to "Actionable" Inference

Most of the data centre growth in the past 24 months has been driven by massive training clusters—thousands of GPUs sitting in hyperscale facilities crunching historical data. But F1 is an inference-heavy sport.

"We are seeing a strategic pivot from the 'Training Era' to the 'Deployment/Inference Era.' As companies like Anthropic move into high-velocity sectors like F1, the infrastructure burden shifts from massive centralised hubs to distributed, high-density nodes."

As AI moves into live sports and manufacturing, the "lag" of a centralised data centre becomes a competitive liability. But you cannot run a real-time race strategy for a car in Monaco from a server farm in Iowa.

We are seeing the early stages of a move toward Edge AI clusters. To support a team like Williams, the compute doesn't just need to be powerful; it needs to be physically or logically closer to the data source. This validates the trend toward regional data centre hubs that can offer sub-10ms latency for critical AI-driven decisions.

Impact on Data Centre Design: Density and Thermal Management

F1 is an exercise in extreme cooling and power density—the exact challenges currently facing the data centre industry. The partnership highlights a looming physical reality: as AI becomes more integrated into real-time operations, the hardware required to run these models is getting hotter and more power-hungry.

- **Rack Density:** We are moving away from the standard 5–10kW racks. AI-specific deployments, particularly those used for complex simulations, are pushing toward 50kW or even 100kW per rack.
- **Liquid Cooling:** Just as an F1 car relies on sophisticated thermal management, the data centres supporting these "thinking partners" are transitioning to liquid-to-chip cooling to maintain performance.

The numbers back this up. Global committed data centre capacity has grown 7x since 2019, and continues to grow strongly, as high-density AI layouts dominate investment commitments. The "F1-ification" of AI will only accelerate the demand for "AI-ready" facilities that can handle the massive power draws required for real-time model interaction.

The Compute Arms Race: Data Centres as the New Engines

Anthropic joining Google (McLaren), Oracle (Red Bull), and IBM (Ferrari) on the F1 grid confirms that compute has replaced oil as the sport's primary fuel. This creates a new "credibility" metric for data centre providers.

If Oracle's cloud helps Max Verstappen win a World Championship, it proves the reliability of their data centres to every enterprise on earth. If Claude helps Williams—a team currently in a rebuilding phase—climb the grid through superior data synthesis, it proves that AI-driven "thinking" is the ultimate competitive advantage. This indicates that Williams aren't just betting on an algorithm; they are betting on the uptime and reliability of the infrastructure behind it.

For data centre operators, the narrative is no longer about being a "utility." It's about being a performance partner. If the AI goes down during a qualifying session, the car stays in the garage. This puts a premium on redundancy and power resilience that exceeds previous industry standards.

Efficiency and the Rise of Small Language Models (SLMs)

There is a common misconception that more AI always means more "everything"—more power, more space, more water. However, the constraints of a racing team (who must travel with their gear and operate within strict cost caps) suggest a move toward efficiency-first AI.

Anthropic's challenge at Williams will likely involve optimizing Claude to run on limited local compute resources. This mirrors a broader trend in the data centre industry: the rise of Small Language Models (SLMs). These are highly distilled, specialized versions of larger models that provide 90% of the utility at a fraction of the compute cost.

For the data centre market, this could lead to a two-tier system:

1. **The Titans:** Massive facilities for training the next generation of Claude.
2. **The Sprinters:** Highly efficient, modular data centres designed specifically to run SLMs for local, real-time applications. ■

AI at the goalposts: Gloucester Rugby's matchday digital transformation

Cheltenham-based technology leader, Emerge Digital, has announced the next chapter in its long-standing relationship with Gloucester Rugby, unveiling a three-year deal that promises to redefine matchday experiences and deepen community engagement for the club.

This latest agreement will extend the partnership to nearly a decade of working together, during which Emerge Digital has helped Gloucester Rugby transform its digital infrastructure, from installing a robust wired network and integrating big screens to relocating staff offices and upgrading team technology. This is all underpinned by ongoing managed IT support, matchday assistance, and cybersecurity solutions.

This next phase will see Emerge Digital and Gloucester Rugby both invest in the club's premier hospitality space, the Captain's Lounge, to boost attendance by 20% this season. Beyond matchdays, it will provide a tech and conferencing venue for local businesses, charities, and community groups.

Upgrades to the lounge include an AI Fan Photo Stand, one of the UK's first real-time AI sport photo experiences, allowing fans to generate high-quality branded images of themselves with the Gloucester Rugby Team.

Future enhancements to the Captain's Lounge will incorporate cutting-edge features designed to elevate the overall fan experience. These improvements include a visually impressive entrance, creating an immersive experience. Additional AI-driven features will continue to roll out throughout the season aimed at enhancing conference and corporate events held in the lounge, setting a new benchmark for smart spaces in rugby.

Nigel Church, CEO of Emerge Digital, which operates across the south west and south Wales, said: "Gloucester Rugby has been a client of Emerge Digital for many years, during which time we've worked closely with the club as its infrastructure and technology have evolved to meet the demands of modern professional rugby. Our partnership has spanned connectivity, application development, and data reporting, supporting both day-to-day operations and initiatives designed to enhance the supporter experience.

"Building on that long-standing relationship, we're proud to sponsor the club's principal conferencing and hospitality space at Kingsholm Stadium. Formerly the Captain's Lounge, the space is now known as the Emerge Digital Captain's Lounge.

"We've already invested significant time, effort, and resource into elevating the experience within the lounge, including the introduction of an AI photo booth that adds a fun, interactive element for supporters on matchdays. Further enhancements are underway, with a continued focus on innovation and using technology to create memorable experiences both on and off the pitch."

Alex Brown, Gloucester Rugby CEO, added: "Our fans want more than a great game - they want convenience, innovation, and experiences that last beyond the final whistle. Partnering with Emerge Digital allows us to deliver that while opening doors for local businesses and community groups. This is about building a future-ready club that leads both on and off the pitch." ■


Nigel Church - CEO of Emerge Digital.

Alex Brown, Gloucester Rugby CEO.

Register for Networking+

Is AI in cybersecurity an opportunity or a risk for MSSPs?

Haris Pylarinos, Founder and CEO of Hack The Box.

AI is rapidly changing how cyber defence is delivered. For managed security service providers (MSSPs), AI-enabled analytics, autonomous monitoring, and automated response offer the ability to analyse more data, close investigations faster and support customers seeking 24/7 assurance, without expanding SOC teams.

AI has the potential to improve cybersecurity, but speed alone doesn't create resilience. If autonomous decision-making is not validated, tuned, and supervised, fast responses may be incorrect. The advantages of automation must be balanced by confidence that it behaves correctly under real-world conditions. Without this, introducing AI into an organisation's cybersecurity becomes a new operational risk.

AI, therefore, represents both a major opportunity and a potential risk for MSSPs, depending on how it is deployed and governed.

Adversaries are already scaling

AI is not only transforming defence, but it has also lowered the barrier for attackers. Generative models now support reconnaissance, target profiling, multilingual social engineering and vulnerability research at scale. Adversaries who once needed time, skill and infrastructure can now generate convincing phishing content or reconnaissance outputs automatically and at volume.

MSSPs cannot defend effectively using static workflows and pipelines. Their response will inevitably require some defensive automation, but this automation must be demonstrably trustworthy. If MSSPs adopt AI tools simply to match attacker speed, without governance and validation, their customers inherit a new risk. One that has been created by the defender, not the attacker.

In this context, AI is not inherently safe or unsafe. Its impact depends entirely on how rigorously MSSPs validate its behaviour.

Benchmarking must become continuous

Historically, SOC processes were validated through staged testing, rule tuning and controlled pilots. That is not enough when dealing with systems that evolve dynamically and make probabilistic decisions. Machine-assisted detection and autonomous containment do not remain static, so static assurance is not enough.

Trust in defensive automation has to be continuously earned. The best way to achieve this is through ongoing benchmarking against realistic adversarial behaviour. Automation should be validated under conditions that replicate genuine threat scenarios. Likewise, analysts must compare how a machine reacts versus how a skilled human would react.

MSSPs need a way to measure human and machine performance in parallel, under pressure and over time. And when adversaries evolve, automation must be challenged again. The aim is continual, evidence-based improvement.

Upskilling is the new strategic priority

AI changes the role of the human analyst; it doesn't eliminate it. Automation delivers the most immediate benefit when it is applied to operational workload, such as triaging repetitive alerts and first-line investigation. When machines handle these tasks, analysts are able to focus on strategic and evaluative responsibilities.

Human analysts also need to know when automated recommendations should be followed, questioned or escalated. They need to interpret ambiguous outputs, recognise model drift and challenge actions that could cause business disruption or legal exposure.

In short, analysts will evolve from process operators into supervisory decision-makers.

This demands new training. Along with attacker techniques and incident response principles, analysts need to learn how to supervise AI reasoning and test machine behaviour under adversarial conditions.

Upskilling is essential for situational awareness and can be a key differentiator for MSSPs. Those that invest in talent capable of supervising and challenging automation will be safer, more mature and more transparent than those that rely on automation without scrutiny.

Upskilling that supports humans and machines together

Security training traditionally focuses on improving human capability. For an AI-powered SOC, training must also support continuous machine improvement. A static model exposed only to historic data will never anticipate emerging tradecraft, while an analyst without experience supervising automation may trust it when they shouldn't.

Future SOC environments will need upskilling that enables humans and machines to improve together. These environments should expose automation to realistic offensive behaviour, while allowing analysts to observe and oversee decision pathways. When both humans and AI are challenged repeatedly, the result is measurable uplift in model precision, analyst judgment and hybrid response quality.

This also reflects a wider principle that security capability is dynamic, so it must be exercised continuously.

The hybrid SOC

The idea of a fully autonomous SOC may sound good in theory, but it isn't realistic. Cybersecurity decisions have commercial, regulatory and reputational consequences, so they cannot be delegated to an AI model without oversight.

The more credible and sustainable operating model is hybrid. Machines accelerate monitoring, triage and early containment; while humans validate, contextualise and govern decisions that have impact. The hybrid SOC is a necessity for resilience in environments where the volume of events and the consequences of misjudgement are high.

For MSSPs, the future will be less about who can automate fastest and more about who can prove their automation behaves predictably over time. Customers will increasingly value evidence that humans and machines are benchmarked against realistic adversarial behaviour, that analysts are trained to supervise automation appropriately and that continuous learning underpins the security posture.

Resilience will not be defined by speed; it will be defined by trust that is created through continuous benchmarking, rigorous upskilling and transparent hybrid oversight. When these safeguards are in place, AI becomes not just safe but strategically valuable. ■

What's on and new...

DATA CENTRE WORLD
4-5 March 2026 Excel London
www.techshowlondon.co.uk

Taking place on 4–5 March 2026 at ExCeL London, Tech Show London brings together five co-located events—Cloud & AI Infrastructure, DevOps Live, Cloud & Cyber Security Expo, Big Data & AI World, and Data Centre World—creating the most comprehensive technology conference programme in the UK.

Tech Show London will once again host its Women & Diversity in Tech initiative — bringing together leaders, advocates, and rising voices for an afternoon of insight, influence, and meaningful connection

The Tech Show London Mainstage Theatre features big names talking about big ideas. This diverse line-up ranges from world-leading experts in AI, Inventors, C-level executives and government representatives, to world renown entrepreneurs, and inspirational voices in tech.

Addressing critical challenges

Designed for senior leaders and practitioners, the 2026 programme addresses the most pressing hurdles for modern organisations:

- Scaling AI Responsibly: Transitioning from "AI hype" to tangible enterprise agents and sovereign AI development.
- Cyber Resilience: Navigating the new 24-hour breach reporting mandates and AI-driven threats.
- Infrastructure Modernisation: Managing the shift to high-density, AI-native data centres and hybrid-cloud efficiency.
- Regulatory Complexity: Understanding the impact of the UK's Data (Use and Access) Act and the Cyber Security & Resilience Bill.

Practical insight across five shows

- Cloud & AI Infrastructure: Hybrid-cloud strategy and the realities of running AI workloads at scale.
- Cloud & Cyber Security Expo: Risk frameworks, ZTNA, and securing modern attack surfaces.
- Big Data & AI World: Data governance, analytics strategy, and the rise of Agentic AI.
- Data Centre World: Next-generation cooling, grid connection solutions, and CNI status implications.
- DevOps Live: Automation, observability, and improving engineering culture.


Tech Show London
Photographer, Aidan Synnott

All conference sessions are free to attend for registered technology professionals and end-users. Attendees will have access to over 300 expert speakers and an exhibition floor featuring more than 400 leading solution providers including AWS, Microsoft, and Google Cloud.

Registration is currently open at www.techshowlondon.co.uk ■

PRODUCTS

At Data Centre World **Technimove** will mark the official launch and showcase of its enhanced Infrastructure Deployment Services, designed to support the next generation of high-performance, AI-driven data centre environments. At the show, Technimove will place the spotlight firmly on its advanced infrastructure deployment capabilities, including AI and GPU-ready deployments, liquid cooling solutions, high-performance cabling, and the end-to-end support services required to design, deploy and maintain these complex environments.

Centiel has announced it will showcase the fault tolerant ride through of its multi-award-winning, true modular, three-phase UPS, StratusPower on stand E110. This year, the Centiel team will be available to demonstrate the sustainable by design, high efficiency Fault Tolerant Ride Through capability of StratusPower and discuss how it's backed by Centiel's 24/7/365 service network and worldwide support. Visitors will have the opportunity to explore the documented and testable resilience of its true-fault tolerant modular architecture, along with its fast delivery times and flexible production capacity.

Mitsubishi Electric is using Data Centre World to launch the ME-CDU which acts as an interface between liquid cooling servers and systems that reuse heat. The unit, which will be on Stand C35 at the Show, is available in a single compact module to provide a seamless integration with hybrid cooling applications and heat reuse or rejection systems, the unit helps ensure a continuous operation of the white space in data centres. It joins the company's comprehensive range of chillers, heat pumps, fan walls, CRACs and CRAHs.

TSL Brings Unified Control and Smarter Power Design to Data Centre World 2026 - Stand: F195. TSL will discuss the potential for bringing together control of multiple, disparate systems in your data centre or other facility into a single, operator-friendly interface. Power management, facility control, and wider systems such as lighting, cooling, HVAC, security, signage, and AV can be monitored and controlled from one consistent Human Machine Interface (HMI), whether deployed on-premises or accessed remotely.

Carrier is expanding its QuantumLeap™ solutions portfolio with the launch of the Carrier AiroVision® 39CV Computer Room Air Handler (CRAH), a high-performance precision cooling solution designed specifically for medium to large-scale data centres. Carrier is part of Carrier Global Corporation (NYSE: CARR), global leader in intelligent climate and energy solutions.

The **Lift&Drive** DC lifters with PRX control system are engineered with advanced ergonomics at their core: lightweight construction, adjustable handlebars, and intuitive joystick controls minimize strain and fatigue. The Patented Ergoswap solution entails a removable battery than can be charged in a safe area that allows you to invest in a more convenient electric lifter.

Frigel, a global specialist in high-efficiency data centre cooling, will present Synergo, its latest integrated solution combining advanced cooling and high-grade heat reuse to support low-carbon, energy-efficient data centre operations. Designed to turn waste heat into a valuable energy resource, Synergo enables seamless integration with district heating networks and industrial heat users, while maintaining resilient, mission-critical cooling performance.

Alfagomma to Present Liquid Cooling Solutions at Data Centre World in London. The company will showcase its advanced liquid cooling solutions specifically engineered to meet the evolving requirements of today's data centres. Designed for Direct Liquid Cooling (DLC) systems, Alfagomma's solutions support In-rack and In-row CDU configurations, ensuring high performance, safety, and operational continuity in both primary and secondary cooling loops.

SPX Cooling Technologies, a full-line, full-service industry leader in the design and manufacture of evaporative cooling towers, fluid coolers and more, has introduced two industrial-grade cooling tower designs for data centres and mission critical locations in Europe, Middle East and Asia (EMEA) markets: the Marley NC Everest® Cooling Tower and the Marley MD Everest® Cooling Tower. These modular, extra-large evaporative cooling towers offer the highest cooling capacity of any factory-assembled cooling towers in the world. (Stand: G165).

A final thought...

When the Natural History Museum set out to redevelop its five-acre gardens, the goal stretched far beyond aesthetic landscaping.

The transformation, made possible through a partnership with Amazon Web Services, was designed to create one of the most data-rich urban nature sites in the UK. What had once been a traditional museum garden is now a digitally enhanced research environment, where soil moisture, insect wingbeats and airborne DNA are monitored with the same enthusiasm as the museum's most prized fossils.

Unveiled as part of the Urban Nature Project, the revamped space invites visitors into a biodiverse landscape that serves a dual purpose: a public green refuge and a fully instrumented scientific observatory capable of running around the clock. The collaboration with AWS underpins this ambition, ensuring that every data point collected in the gardens can feed into ongoing research into how urban ecosystems behave in a rapidly changing climate.

A sensor-first landscape

To bring this idea to life, the museum embedded a network of environmental and acoustic sensors throughout the gardens. These devices continually track temperature fluctuations, humidity levels, water conditions in the ponds, and even the subtle chorus of urban and natural sounds. The underwater microphones capture audio frequencies that the human ear misses, revealing everything from amphibian activity to the acoustic signatures of passing insects. Over time, this system is expected to generate many terabytes of recordings, giving scientists a depth of ecological insight that would be impossible through manual surveying alone.

Alongside the hardware sits a strategy for environmental DNA monitoring. By collecting tiny traces of genetic material from soil and water, researchers can detect organisms that never reveal themselves to the naked eye. This adds a sensitive biological layer to the sensor data, enabling the museum to build a species-level picture of biodiversity that is continually updated. The technology turns what would once have been a static wildlife snapshot into a constantly refreshing view of the living world just outside the museum's doors.

"We're incredibly excited about the switch-on of our new environmental sensor network. This marks a major milestone in transforming our gardens into a living laboratory, helping us better understand how urban nature is changing in real time. With the support of Amazon Web Services, we're now able to collect and share data at an unprecedented scale, deepening our understanding of biodiversity and drive forward science-led nature recovery in the UK's urban spaces," says Ed Baker, Acoustic Biology Researcher at the Natural History Museum.

Cloud infrastructure that binds the ecosystem together

The flow of information only becomes useful once it is processed, analysed and shared, and this is where AWS's cloud services form the backbone of the project. The museum's Data Ecosystem platform acts as a central environment for all incoming data streams, integrating acoustic recordings, eDNA results, microclimate readings and visual observations into a single, scalable system. Built on serverless and managed AWS technologies, the platform allows the museum's scientists to handle large, fast-growing datasets without maintaining physical servers or complicated infrastructure behind the scenes.

The cloud-based environment also makes collaboration far simpler. Researchers based at the museum and across the UK can securely access data and analytical tools, compare findings, and run models that show how biodiversity responds to seasonal change, extreme weather or urban activity. Historical records that date back almost three decades have already been absorbed into the system, allowing today's observations to be set against long-term baselines. Citizen science plays a role too: data from nationwide programmes flows into the same platform, giving non-scientists a meaningful stake in the research and expanding the scale of observation far beyond what museum staff could achieve on their own.

A blueprint for urban biodiversity research

The result of this collaboration is a garden that behaves far more like an outdoor laboratory than a landscaped attraction. Instead of offering occasional glimpses into the species that pass through, the space delivers a continuous, highly detailed narrative of how urban nature shifts from hour to hour and season to season. The project establishes a model that other towns, cities and institutions can adapt, showing how cloud technology, environmental sensors and modern genetic techniques can turn even small urban green spaces into valuable ecological observatories.

"We are delighted to support the Natural History Museum to transform and accelerate its scientific research and community science capabilities with the cloud. By building the Data Ecosystem using cloud technology the Museum's scientists can securely store and process data from the gardens for the first time. This allows the Museum to turn this data into actional insights to support the UK's urban nature recovery. The scientists can continue deepening their understanding of the UK's urban diversity by using the cloud to scale-up the Data Ecosystem as more data from the gardens is collected over time," says Hilary Tam, Sustainability Leader, Europe Middle East & Africa at Amazon Web Services.

As the data continues to accumulate, the museum now has the tools to explore questions that were previously inaccessible: how microclimates evolve within metres of each other, how species respond to extreme temperatures, how urban noise influences wildlife behaviour, and how biodiversity ebbs and flows in real time. The project signals a shift in how museums engage with science, positioning the Natural History Museum not just as a guardian of past life, but as a dynamic centre for understanding the living world that surrounds it today. ■

Ed Baker, Acoustic Biology Researcher at the Natural History Museum and Hilary Tam, Sustainability Leader, Europe Middle East & Africa at Amazon Web Services install an environmental sensor in the pond to measure temperature at various depths. © The Trustees of the Natural History Museum, London.