JUNE 2025
networkingplus.co.uk

NETWORKING+

Transforming the SOC

Staying ahead of modern threats

Zeki Turedi, CrowdStrike, p6

AI and the neural edge

The data centre is no longer a warehouse for computing

Peter Wilcock, Latos Data Centres, p14

Questions and answers

I look up to people who inspire curiosity

Dominic Mensah, Lakeside Software, p16

Industry warns against 'crying wolf' over data breaches

Late in June, many of us witnessed the breaking news story: 16 billion credentials leaked online in 'the mother of all data breaches.'

Researchers from an online cybersecurity news site said that they had found 30 datasets full of credentials harvested from infostealers and leaks. The datasets were exposed only temporarily through unsecured Elasticsearch or object storage instances.

However, multiple incident response specialists, researchers and cybersecurity experts have since disputed those claims and questioned the data and analysis the assertion was based upon. Moreover, industry experts have reported that by comparing released sample data against previous credential leaks, it becomes clear that most of these credentials were from previously released password dumps.

Accordingly, strong warnings have emerged from industry that misinformation or embellishment can be a disservice, drawing attention away from verified attacks and new weaknesses as they are exposed.

Kev Eley, VP UK&I at Exabeam, warns that "the recent hype around the '16 billion password breach' is a prime example of how exaggerated cybersecurity narratives can do more harm than good. Not only is there little evidence to support the claim, but repeated focus around massive leaks can also cause companies to view attacks as noise rather than real threats. When there's such a big focus on these large numbers, attention shifts away from real attacks that cause real damage to businesses – like phishing or ransomware threats. To break this cycle, organisations need to build response plans that rely on facts, not fear."

"A majority of the data in the '16 billion password breach' appear to be old or previously disclosed compromises, causing some to call the various media reports nothing more than fear mongering," agrees Mike Puglia, General Manager – Security Products, Kaseya. However, "while the headlines may be fantastical, we should not underestimate the immense value of aggregating thousands of siloed breaches into a single set of data and having that information exposed to the entirety of the malicious actor community in one release. Even if the majority of passwords are no longer valid, on average over one million will still 'work'. Perhaps more importantly, it gives an extremely large data set for attackers to leverage AI enhanced phishing emails with extreme personalisation."

As such, James Shank, Director of Threat Operations at Expel, says that "if this news frightens you, then your security program probably has some fundamental gaps. Let this be the fuel you need to position yourself and your department for solving the problem systematically, rather than defending against the news du jour. There will always be another breach, with even more passwords, and emergency handling will continue if you don't have systematic defenses in place."

In the wake of the incident, experts are keen to highlight the pervasiveness of infostealer malware, as well as how enterprises should protect against this type of attack. The fact that someone was able to compile 16 billion records from old – sometimes years old – breaches, shows how big the problem is.

"We've seen recent attacks where sensitive, regulated data was stolen — raising serious concerns about fines and loss of trust. It's not just about having backups anymore," notes Charles Burger, Director of Cybersecurity Solutions, Nexsan. "Organisations need immutable storage with detailed access logs to really protect their most valuable data. If credentials are compromised, being able to see exactly what was accessed, by whom, and from where is critical to respond quickly and limit the damage." ■

IN DEPTH: Network scaling p7-8

Project Reach launches huge upgrade to Britain’s rail telecoms

A groundbreaking partnership between Network Rail, Neos Networks, and Freshwave has been launched to address the worst signal blackspots along Britain’s major rail routes.

This collaborative effort marks the most significant enhancement to the country’s rail telecommunications infrastructure in recent decades, promising improved connectivity for millions of passengers and laying the foundation for a new high-capacity telecoms network that will support the future ambitions of Great British Railways and the nation’s broader digital economy.

“This partnership marks a significant step forward in improving the UK’s rail infrastructure. By enhancing connectivity and addressing signal blackspots, we are ensuring that passengers experience a more reliable and efficient service. This initiative not only benefits rail users but also supports our broader goals of economic growth and digital innovation,” said Heidi Alexander, Secretary of State for Transport.

Initially, Neos Networks will deploy 1,000km of ultrafast fibre optic cable along key routes including the East Coast Main Line from King’s Cross to Newcastle, the Chiltern Main Line, the West Coast Main Line to Manchester, and the Great Western Main Line from London to Cardiff. There is an ambition to expand this network beyond 5,000km in the future.

“I’m delighted that we have now signed this innovative deal with our partners Neos Networks and Freshwave. This investment model will deliver the necessary upgrades to our telecoms infrastructure faster whilst offering significant value-for-money for the taxpayer and stimulating wider economic benefits across the country,” said Jeremy Westlake, Network Rail’s chief financial officer. “As we move towards becoming a unified railway with the formation of Great British Railways, the enhanced telecoms infrastructure will play a key role in our ambition to provide a data-driven railway of the future, delivering better connectivity and a better, more reliable train service for our passengers.”

Freshwave will initially focus on installing mobile infrastructure to eliminate signal blackspots in 57 tunnels and deep cuttings along the same major lines, including notable tunnels such as Chipping Sodbury near Bristol and Gasworks and Copenhagen tunnels outside King’s Cross. The company will also work with mobile network operators to upgrade 12 key stations with new 4G and 5G infrastructure, with the first installations expected to commence in 2026.

The upgraded network, which will expand from 48 to 432 high-count fibre cables, will enable Network Rail to monitor assets more effectively and support the deployment of new technologies such as trackside sensors and CCTV, enhancing safety and operational efficiency. The project’s innovative commercial approach combines public and private investment, expected to save taxpayers around £300 million while delivering substantial benefits to rail users.

Neos Networks, supported by Infracapital and SSE, will deploy high-capacity fibre along the tracks, utilising spare fibre capacity to upgrade critical telecoms infrastructure and create a resilient digital backbone supporting the UK’s broader digital ambitions. Lee Myall, CEO of Neos Networks, highlighted the company’s commitment to supporting Britain’s critical infrastructure and fostering innovation through enhanced connectivity.

This ambitious project exemplifies a collaborative approach to modernising Britain’s rail infrastructure, bringing together public and private sectors to deliver a smarter, more connected railway for the future. ■

Ark Data Centres partners with Nebius to launch UK’s first NVIDIA Blackwell Ultra GPU Cluster

Ark Data Centres has announced a strategic, long-term partnership with Nebius, a prominent AI infrastructure company. This collaboration will see the installation of one of the UK’s first deployments of NVIDIA Blackwell Ultra GPUs at Ark’s Longcross Park campus in Surrey.

The initial rollout will feature 4,000 of these cutting-edge GPUs, creating a powerful AI compute cluster aimed at supporting UK start-ups, research institutions, enterprises, and public-sector bodies — including the NHS — to develop and scale artificial intelligence using the most advanced computing technology available.

The GPU deployment aligns closely with the UK government’s AI Opportunities Action Plan, which aims to expand the nation’s domestic AI compute capacity. Beyond bolstering the capabilities of UK AI developers, the infrastructure supplied by Nebius is expected to stimulate job creation and attract additional investment into the country’s burgeoning AI sector.

Under the terms of the agreement, Nebius will occupy purpose-built, liquid-ready data hall capacity at Longcross Park, a facility designed from the ground up for high-density AI workloads. The campus offers the high power density, advanced cooling systems, low-latency connectivity, and sustainability features necessary to support next-generation accelerators, ensuring the infrastructure can meet the demands of rapidly evolving AI applications.

“Partnering with Ark provides us with access to an environment built specifically for the next wave of GPU innovation, equipped with the infrastructure needed to support today’s highly intensive AI cloud workloads. By locating compute resources close to the UK’s leading start-ups, enterprises, researchers, and public-sector innovators, we can help them move faster from initial ideas to tangible implementation,” said Andrey Korolenko, Chief Product and Infrastructure Officer at Nebius.

This new GPU cluster will showcase Ark’s ability to support demanding AI workloads, leveraging high-efficiency cooling, resilient power systems, and dedicated on-site power generation—capabilities that many legacy data centres struggle to deliver at scale.

“AI clouds require sites capable of delivering high-density cooling today and scaling rapidly for tomorrow’s demands. Longcross Park was engineered precisely for that purpose, and we are proud to welcome Nebius as a key AI tenant in the UK. This partnership demonstrates that the UK already has the necessary infrastructure to support global AI innovation and underscores our commitment to investing £7.5 billion in new, AI-ready capacity across the country to underpin the UK’s AI future,” said Huw Owen, CEO of Ark Data Centres.

Nebius’s first investment in the UK marks a significant milestone in its global buildout of AI infrastructure. With this deployment, Nebius will operate seven AI clusters across six countries in Europe, the US, and the Middle East, positioning it as one of the largest independent AI infrastructure providers worldwide.

The company’s deployment of thousands of NVIDIA Blackwell Ultra GPUs in the UK is expected to be operational by Q4 2025, enabling local AI innovation at scale and further cementing the UK’s position as a global hub for AI development. ■

Half of enterprise Windows endpoints have migrated to Windows 11 ahead of support deadline

ControlUp has released new insights from its Windows 11 Readiness report, revealing that 50% of enterprise Windows endpoints have now completed their migration to Windows 11.

This marks notable progress from last year, when over 82% of enterprise devices were still running Windows 10. With less than four months remaining before Microsoft’s support for Windows 10 ends on 14 October 2025, the data underscores both encouraging advancements and ongoing challenges in enterprise migration efforts across industries, regions, and organisation sizes.

“While reaching 50% migration is a significant milestone, organisations shouldn’t relax,” said Marcel Calef, America’s Field CTO at ControlUp. “With Windows 10 support ending soon, companies need to accelerate their plans to avoid being caught unprepared. Our data shows that migration progress is uneven, and many enterprises still face substantial hardware and planning hurdles.”

Analysis of over one million enterprise endpoints shows that the Education and Technology sectors are leading the migration, with 77% and 73% of their devices already upgraded to Windows 11, respectively. Conversely, the Healthcare sector lags behind at 41%, and Finance at 45%. Further inspection indicates that 19% of Healthcare endpoints require complete hardware replacement before supporting Windows 11, compared to just 3% in Finance.

The report also highlights that organisations with over 10,000 Windows devices are the least prepared — only 42% have completed their migration. The complexity of large IT environments and the prevalence of legacy hardware contribute to these hurdles, emphasising the importance of early assessment and planning. ■

EDITORIAL:

Editor: Amy Saunders

Designer: Ian Curtis

Sub-editor: Gerry Moynihan

Contributors: Charlie McRae, Zeki Turedi, Paul Speciale, Peer Wilcock, Matthew Terry, Dominic Mensah

ADVERTISING & PRODUCTION:

Sales: Kathy Moynihan
kathym@kadiumpublishing.com

Production: Karen Bailey
karenb@kadiumpublishing.com

Publishing director:
Kathy Moynihan
kathym@kadiumpublishing.com

Networking+ is published monthly by:
Kadium Ltd, Image Court, IC113, 328/334
Molesey Road, Hersham, Surrey, KT12 3LT

Tel: +44 (0) 1932 886 537

ISSN: 2052–7373

Outdated warehouse technology a barrier to business growth

Research from Inteq reveals that the majority of UK retailers feel constrained by aging warehouse and fulfilment systems, which are hampering their ability to expand and respond to fluctuating customer demand.

A survey of retail and eCommerce senior leaders found that 58% believe their current fulfilment infrastructure limits their growth prospects, while 56% say warehouse logistics issues are directly affecting their business expansion.

The study highlights the broader challenges faced by the sector amid shifting consumer expectations, ongoing supply chain disruptions, and seasonal peaks. Retailers recognise that outdated facilities and systems are a significant obstacle to maintaining operational performance and delivering high-quality customer service.

A striking 88% of respondents agree that integrating robotics and automation into their fulfilment processes would give them a competitive edge during demand surges, such as holiday shopping seasons and major sales events. Many retailers are already experiencing the benefits of such investments: nearly 70% report tangible improvements within less than a year of implementation. The key advantages include increased operational efficiency (41%), greater scalability (28%), and enhanced accuracy in order processing (26%).

Given the rapid pace of technological change, retail leaders see modernising warehouse logistics as essential for staying competitive. They are prioritising upgrades to their infrastructure and technology, driven by the need to meet the rising demand for fast, error-free deliveries.

However, the process of adopting new systems presents challenges. The study found that selecting the right technology partner is critical for success. Ongoing system maintenance concerns were raised by 34% of respondents, while 32% cited difficulties with technical integration. Retailers emphasise the importance of careful partner evaluation to prevent being locked into inflexible systems that may hinder future automation or growth initiatives.

Despite hurdles, most retailers remain committed to investing in automation to address demand volatility and enhance customer experience. The research indicates a swift return on investment, with 69% of companies experiencing efficiency gains within twelve months. The findings are further explored in the ‘Fit for Growth’ report, which offers strategic insights on balancing current demand with future expansion, emphasising how robotics and automation can drive resilience and profitability in a highly competitive market. ■

The Walton Centre transforms clinical operations

The Walton Centre, a specialist neurosciences hospital in the UK, has successfully implemented a comprehensive digitisation of its patient records, resulting in notable improvements in clinical efficiency, patient safety, and staff wellbeing.

This initiative was launched in response to a 2019 Care Quality Commission (CQC) inspection, which highlighted concerns over the hospital’s mixed use of paper and digital notes, posing potential clinical risks. To address this, the Trust undertook a competitive tender process that led to selecting Apogee’s document scanning and data capture services.

Today, the majority of patient records have been scanned and integrated into the hospital’s electronic patient record (EPR) system. This integration allows clinicians instant access to complete, real-time patient information, including medical histories, treatment plans, and device data, directly at the point of care. The project also includes an in-house scanning facility that ensures newly generated patient documents are added to the EPR within 48 hours of outpatient visits or within seven days following hospital discharge.

The impact of this digitisation effort has been profound. It has enabled the hospital to redeploy its medical records staff to higher-value tasks, thereby improving patient flow and operational efficiency. One of the most significant achievements has been clearing a backlog of 5,000 patient referrals. Additionally, new patient registrations are now conducted upon arrival, streamlining administrative processes. The move away from paper records aligns with the hospital’s broader goal of becoming ‘paper light’ by 2026 and achieving digital maturity in line with NHS Digital Strategy.

“They are more than just a supplier — they are an integrated extension of our team. We were initially nervous about digitising such a high volume of records, but Apogee responded with professionalism, expertise, and valuable solutions, helping us meet our operational goals,” said Sam Holman, Head of Transformation and Corporate Operations at The Walton Centre. ■

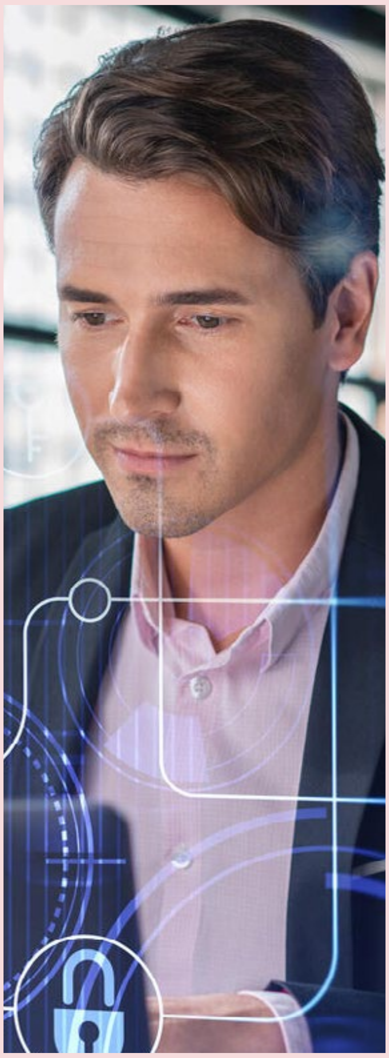

Research highlights growing cloud security risks

The Tenable Research 2025 Cloud Security Risk Report reveals critical gaps in safeguarding sensitive data, managing identities, and securing cloud workloads — especially with the increasing use of AI resources.

The report highlights concerning levels of data exposure, noting that 9% of publicly accessible cloud storage contains sensitive information, with 97% of this data classified as restricted or confidential. Such exposure significantly raises the risk of exploitation by malicious actors, particularly when misconfigurations or embedded secrets like passwords and API keys are present. Security inconsistencies across major public cloud providers — Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure — compound these vulnerabilities, leaving organisations exposed to potential breaches.

In terms of secrets and workload security, the report reveals that more than half of organisations (54%) store at least one secret directly within AWS Elastic Container Service (ECS) task definitions, creating a direct attack vector. Similar patterns are evident on GCP Cloud Run, where 52% of organisations store secrets within resources, and 31% of Microsoft Azure Logic Apps workflows contain embedded secrets. Additionally, 3.5% of all AWS EC2 instances were found to hold secrets within user data, representing a substantial security concern given AWS's widespread adoption.

While there has been some progress — specifically, a reduction in the 'toxic cloud trilogy' scenario, where workloads are publicly exposed, vulnerable, and highly privileged — from 38% to 29% — the risk remains significant. This persistent threat underscores the importance of tightening security controls around cloud workloads.

The report also emphasises ongoing challenges in identity and access management. Although 83% of AWS organisations use Identity Provider (IdP) services to manage cloud identities — a recognised best practice — risks persist due to overly permissive default settings, excessive permissions, and lingering standing privileges. These misconfigurations open pathways for attackers to exploit over-privileged accounts, access sensitive assets, or extract embedded secrets with relative ease.

"Despite the numerous security incidents we've observed, many organisations continue to leave critical cloud assets exposed through avoidable misconfigurations. Attackers can exploit public access, embedded secrets, or overprivileged identities to gain entry," said Ari Eitan, Director of Cloud Security Research at Tenable.

To combat these vulnerabilities, Eitan advocates for continuous, proactive risk management. He emphasises that security teams require comprehensive visibility across their cloud environments and the ability to automate remediation to prevent threats from escalating. The report advocates for a unified approach to cloud exposure management, increased asset visibility, and systematic automation of security processes — especially as reliance on AI-driven cloud resources continues to grow. ■

Legacy kit threat outlined

Essenkay has issued a stark warning to UK companies relying on legacy technology, suggesting that such dependence could be exposing them to serious security threats and operational inefficiencies.

The consultancy highlights increasing concerns across various sectors — including manufacturing, engineering, logistics, and professional services — that continue to depend on aging systems despite the availability of advanced cloud solutions.

A recent study by the Financial Conduct Authority (FCA) underscores the extent of the problem within the financial sector, revealing that 92% of financial firms still rely on legacy technology, with 78% of their data stored on on-premise systems. Essenkay emphasises that this challenge is widespread and particularly pressing amid mounting economic pressures on UK businesses.

While legacy platforms may seem to function adequately, they quietly hinder business productivity, create operational blind spots, and introduce significant cybersecurity vulnerabilities. As economic conditions become more challenging, maintaining such systems is not only inefficient but increasingly risky.

Operational inefficiencies are a primary concern. In manufacturing, outdated systems can lead to inaccuracies in stock records, while in logistics, fragmented supplier data hampers operational flow. These legacy platforms often force employees into manual workarounds, which are time-consuming, error-prone, and reduce overall efficiency.

Another critical issue is the proliferation of data silos. Without a modern Enterprise Resource Planning (ERP) system, vital business information can become dispersed across spreadsheets, local servers, and obsolete software, impairing reporting and decision-making processes.

Cybersecurity remains a growing concern, especially as cyberattacks increase in frequency. Government data indicates that half of UK businesses suffer security breaches annually, with recent incidents at major retailers like Marks & Spencer highlighting vulnerabilities. Outdated devices such as old barcode scanners, industrial IoT components, and legacy financial platforms are potential weak points.

Beyond operational risks, Essenkay warns that reliance on obsolete systems can cause companies to miss vital business opportunities. Limited access to advanced analytics and predictive tools prevents organisations from identifying customer trends or operational efficiencies. ■

Investigation highlights rising cyber threats amid Iran conflict

A recent investigation by Sky News has revealed an increase in cyberattacks linked to the Iran conflict, targeting organisations across various sectors.

In response, UK Prime Minister Sir Keir Starmer, speaking at the NATO Summit, called on businesses of all sizes and industries to prioritise cybersecurity and immediately review and bolster their defenses.

While the warning underscores the urgency of the threat, many organisations are already experiencing politically motivated cyberattacks, highlighting the need for heightened vigilance.

Clinton Groome, CEO of Espria, emphasised the importance of proactive measures:

"As tensions escalate globally, threat actors will continue to exploit digital vulnerabilities, and even neutral businesses risk being caught in the crossfire. This situation underscores a crucial lesson—businesses should not wait for government alerts to act. Instead, IT leaders need to invest in integrated defenses, educate their users, and develop a clear cybersecurity strategy now."

A persistent challenge for organisations remains human error, which cyber experts warn is often exploited by attackers. Groome advocates strengthening the human firewall—employees—by fostering cybersecurity awareness across the entire organisation.

"Cyber resilience isn't just about technology; it's about people. Social engineering attacks thrive on distraction, fear, and information overload," Groome explained. "Shockingly, a BT study found that 39% of SMEs — around 2 million businesses — haven't even provided cybersecurity training for their staff, leaving them dangerously vulnerable."

He recommends regular training exercises, incident response drills, and ongoing reinforcement of best practices to cultivate a security-conscious workforce. Combining these efforts with layered defenses — such as multi-factor authentication, timely patching, and securing IoT devices — can significantly reduce human-related risks and limit potential fallout. ■

Word on the web...

AI-washing: the risks of jumping on the AI bandwagon

James Clark, Data Protection, AI and Digital Regulation, partner at Spencer West LLP

To read this and other opinions from industry luminaries,

visit www.networkingplus.co.uk

Real Security for the Real World

Real Security for the Real World
Explore how Real Security from WatchGuard can help your business

Cybersecurity is hard. But solutions that can't handle today's real, imperfect operating conditions make it so much harder. That's where WatchGuard is different. With WatchGuard, you get Real Security for the Real World. We focus on solving your real problems – the everyday obstacles standing between you and effective, scalable, reliable, and economical security.

ENDPOINT SECURITY

IDENTITY SECURITY

NETWORK SECURITY

MANAGED SECURITY

Tel: +44 (0) 125 661 0130   uksales@watchguard.com   www.watchguard.com

SOC transformation: the key to staying ahead of modern threats

Zeki Turedi, Field CTO, EMEA, CrowdStrike

Cyber adversaries continuously pose new and highly sophisticated threats that evolve beyond the capabilities of legacy security solutions — a fact that IT decision-makers cannot afford to ignore. Modern attacks demand advanced solutions, yet traditional tools available to security operations centre (SOC) teams are increasingly ill-equipped to meet these demands.

This reality underscores an urgent need to adapt and future-proof legacy security information and event management (SIEM) tools. By leveraging the latest technologies, SOC teams can improve and streamline their operations. However, such a transformation takes time — which ultimately benefits threat actors.

Security decision-makers must recognise the growing capabilities of today's adversaries and prioritise the modernisation of their organisation's SOCs. Starting this necessary transition is the first step toward ensuring the integrity of organisational defences in 2025 and beyond.

SOC teams are being left behind

Today's adversaries can infiltrate organisations at unprecedented speeds: the fastest recorded eCrime breakout in 2024 was just 51 seconds. But it's not just speed that makes modern threats more dangerous. Attackers are becoming stealthier and increasingly imaginative in their methods. The stark reality is that as adversaries evolve, SOC teams risk being left behind.

Holding security teams back are legacy SIEMs, plagued by poor scalability and slow, manual security investigations. These outdated solutions diminish SOC effectiveness by flooding analysts with excessive, irrelevant, or duplicate data, making it nearly impossible to cut through the noise and effectively respond to threats.

Adding to the challenge, SOC teams struggle with the exponential increase in data volumes. Legacy SIEMs, with their outdated billing models based on ingest data volume, force security teams into a difficult trade-off: prioritise budget or security. This economic constraint limits their ability to log and retain critical data, creating security blind spots that adversaries can exploit.

A connected approach

If SOC teams are to keep pace with today's adversaries, they need tools that match the speed, scale, and intelligence of modern threats — not outdated solutions that hinder their ability to detect and respond effectively. The answer is a next-gen SIEM that takes a connected approach to security operations, converging data, AI, and workflow automation into a unified cybersecurity platform.

Next-gen SIEMs address many of today's unique SOC challenges, particularly data ingestion and storage, by having critical security data built in from the start as part of a unified security platform. As a result, security teams no longer need to spend countless cycles on data onboarding, allowing them to operationalise security insights immediately.

If data is the heart of SIEM, then detection content is the brain. With AI-powered detections, next-gen SIEMs are smarter than their legacy counterparts, enriching data and correlating it with comprehensive threat intelligence and security telemetry from endpoints, identities, workloads, and more. By putting data into a risk and security context, analysts receive higher-fidelity, more actionable alerts, cutting through noise and eliminating the endless false positives that plague traditional SIEMs.

Unleashing the power of the SOC team

SOC teams are already reaping the benefits of streamlined security operations by adopting modern technologies like next-gen SIEMs. These advancements help security professionals cut through digital noise, reduce false positives, and accelerate investigations, leading to faster and more precise threat detection.

Across the SOC, employees working with next-gen SIEM experience significant gains in efficiency. Security analysts can leverage automated workflows to streamline incident response and analyse threats at unprecedented speed. Meanwhile, with rapid search speeds compared to legacy SIEMs, threat hunters can search for adversaries and incidents faster, allowing them to proactively uncover threats before they escalate.

Unlike traditional SIEMs, which rely on manual processes and static rules, next-gen SIEM's AI-native approach continuously adapts to evolving threat patterns, uncovering sophisticated attacks that might otherwise go undetected. Integrating a next-gen SIEM empowers security teams to shut down adversaries at unrivaled speeds while simultaneously reducing overall SOC costs.

Beyond operational efficiency, implementing these technologies has been shown to significantly reduce security costs, namely data ingest — a crucial factor as many SOC teams operate under tight budgets and limited staffing. By embracing next-gen SIEM, SOCs can achieve more with fewer resources, strengthening security without increasing operational burden.

SOC transformation: A non-negotiable imperative

As attackers grow faster and employ sophisticated tactics, including malware-free and identity-based attacks, fragmented security tools are becoming a liability. These disjointed systems create blind spots and waste valuable time, a luxury that SOC analysts can't afford in the face of modern threats.

To keep pace, SOC teams must embrace more agile, integrated solutions that enhance efficiency and accelerate response. While transitioning to a modern security architecture takes time, the benefits far outweigh the effort. Moving forward, adopting a modern, AI-driven SOC approach will be essential to stay ahead of evolving threats and ensure operational resilience. ■

Ransomware Is Rising—Is Your Backup Strategy Ready?

In an increasingly digital world, data backups are often treated as standard IT procedures. Yet, their importance becomes critical when cyber threats strike. With ransomware attacks rising at an alarming pace, Kingston Technology emphasises the necessity of the 3-2-1 backup strategy as a robust and proactive defence for safeguarding both enterprise and personal data.

The Ransomware Wake-Up Call

Recent trends highlight the urgency of this approach. According to the UK Government's 2025 Cyber Security Breaches Survey, the proportion of UK businesses reporting ransomware incidents has doubled—growing from under 0.5% in 2024 to approximately 1% in 2025. That's an estimated 19,000 affected organisations.

The retail sector has been especially impacted. Between April and May 2025, high-profile UK retailers—Marks & Spencer, Co-op, and Harrods—fell victim to coordinated ransomware attacks, resulting in major operational disruptions and financial damages in the hundreds of millions of pounds.

Once an organisation falls victim to ransomware, recovery of encrypted data is rarely possible—even if the ransom is paid. The 3-2-1 backup strategy offers a significantly more reliable and cost-effective means of safeguarding data.

A Smarter Strategy for Resilience

The 3-2-1 method is a robust, easy-to-follow strategy that ensures data resilience. It involves keeping three copies of your data—the original plus two backups—stored across two different types of storage media to protect against hardware failure or malware. Crucially, at least one backup is kept off-site and isolated from the network, protecting against physical disasters and ransomware attacks. This scalable framework helps organisations maintain business continuity and recover swiftly from data loss.

To enhance resilience, Kingston recommends using encrypted hardware for offline storage. Devices such as the Kingston IronKey VP80ES SSD and IronKey VP50 USB drive are built to withstand brute-force attacks and prevent unauthorised access—adding another critical layer of protection.

Implementing the 3-2-1 strategy significantly reduces the risk of irreversible data loss and enables faster recovery—all at a fraction of the cost of a ransomware breach.

Scale or stumble: the enterprise network crossroads

Strategic scalability, not just capacity, is defining the next generation of enterprise networks. Here's how tech leaders are navigating complexity with flexibility, automation — and a sharp eye on ROI.

Today's enterprise networks aren't just under pressure — they're under siege. Exploding traffic from cloud services, AI workloads, IoT expansion, and hybrid work models has made static infrastructure a liability. In response, IT leaders are rethinking their architectures from the ground up, championing agility, automation, and intelligent design to future-proof their networks.

From rigid to resilient: the architecture of sustainable scale

In the face of escalating demands — from AI workloads and IoT proliferation to hybrid workforces — traditional network designs are fast becoming obsolete. What enterprises need now is adaptability.

Mark Burski, Managing Director at Digital Carbon, reports that agility must be baked into the architecture: "to enable sustainable scaling as enterprise network demands grow, I recommend architectures that combine intelligent automation, application awareness, and integrated security."

"Enterprises today face relentless growth in user devices, cloud services, and security requirements; all of which place massive strain on traditional three-tier network designs. To scale sustainably, enterprises need to adopt a modular, cloud-aligned architecture that's grounded in Zero Trust principles and built on technologies like Cisco SD-Access in the campus, paired with Secure SD-WAN at the edge," asserts Hamzah Malik, Presales Solutions Consultant, CACI Ltd.

Software-defined networking, especially SD-WAN, plays a key role by dynamically optimising traffic across diverse connections-such as fibre, 5G, and satellite-ensuring critical applications maintain performance as networks expand. Incorporating AI-driven analytics further allows networks to anticipate and adapt to changing traffic patterns, automatically prioritising bandwidth for essential workloads.

"These technologies have enabled organisations to maintain reliable

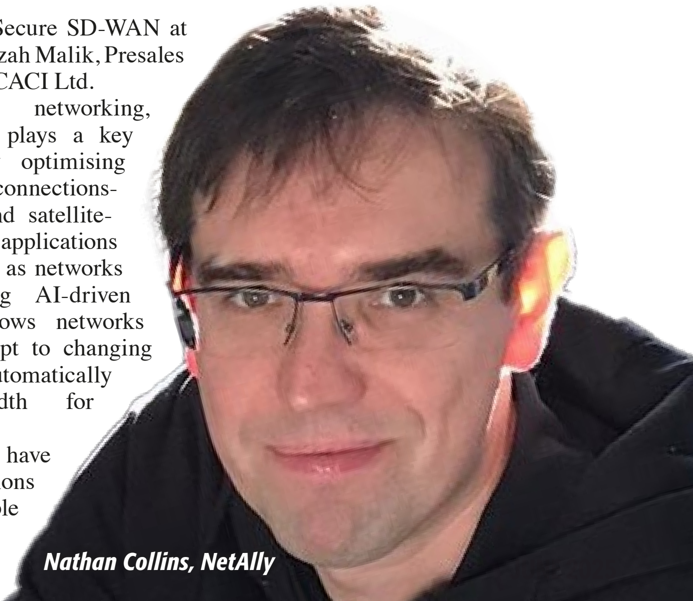

Nathan Collins, NetAlly

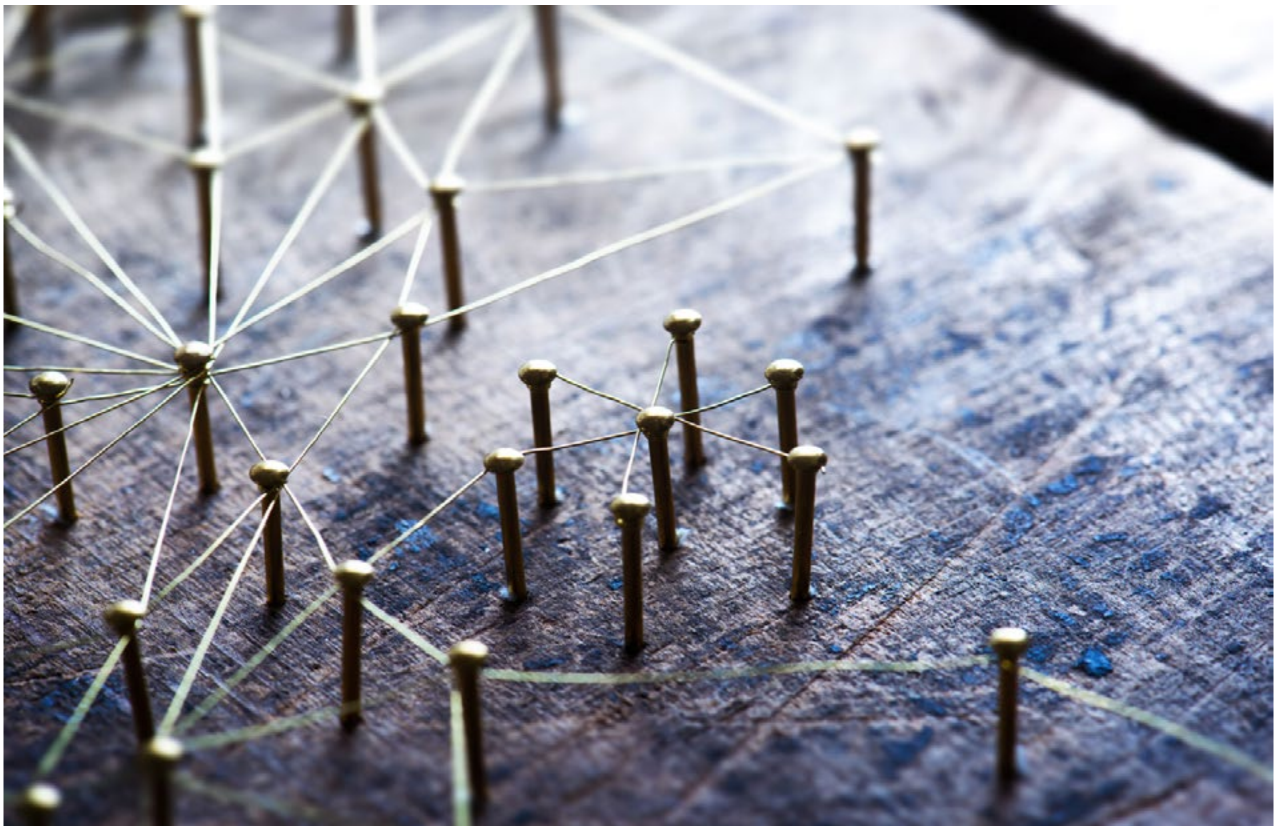

connectivity and application performance across distributed sites, including those using variable wireless links," reports Burski.

"As enterprise networks absorb the weight of digital transformation, IoT expansion, and hybrid work, the emphasis must shift from static capacity to adaptive scalability. Architectures that blend Wi-Fi upgrades, PoE+ infrastructure, and edge-aware visibility are essential to supporting both bandwidth and resilience," agrees Nathan Collins, VP EMEA at NetAlly. "Organisations that delayed Wi-Fi upgrades to extend asset life are now facing diminishing returns — performance bottlenecks, power mismatches, and rising support costs. Sustainable scaling starts with a clear baseline: site-level audits, spectrum scans, and PoE verification provide the actionable insight required to deploy high-density tri-band access points reliably."

"As enterprises face mounting pressures to scale their networks, it's no longer optional to adopt flexible, scalable architectures like SD-WAN, cloud-based networking, and software-defined networking (SDN)," notes Michael Hern, Head of Networks Engineering, razorblue. "These are no longer just 'nice-to-haves' – they are essential to managing growing demands. Technologies like these not only enable more efficient bandwidth management and greater agility but also bolster security, which is critical in an era of increasing cyber threats."

Dynamic by default

With the bursty nature of modern workloads, forecasting network needs has become less of a spreadsheet exercise and more of an intelligent, ongoing process.

"Enterprises should approach network capacity planning as a dynamic, ongoing process rather than a one-off exercise," says Burski. "With the rise of AI, cloud, and distributed applications, traffic patterns have become more unpredictable and bursty, often requiring high bandwidth in both directions and low latency for critical workloads. Relying solely on historical usage or static forecasts is no longer sufficient."

"Best practice is to review these forecasts quarterly, correlating them with business events (product launches, sales promotions) and updating your network as code templates accordingly. By treating your network as a living system where capacity is forecast, tested, and adjusted in small increments, you stay agile and capital efficient, even as demands evolve," adds Malik.

Indeed, modern capacity planning should leverage real-time analytics and intelligent automation to monitor application behaviours and network performance continuously.

"Using software-defined networking technologies, enterprises can dynamically allocate resources, prioritise essential applications, and adapt to sudden surges in demand. Automated traffic steering and dynamic path optimisation help ensure that critical services maintain performance, even during unexpected spikes," adds Burski.

As enterprise networks grow, the manual workload threatens to outpace the people managing it.

"Manual device changes are a recipe for delays, configuration drift, and human error. By codifying your network in Infrastructure as Code, you gain repeatability, version control, and instant rollback," notes Malik. "Combine that with AI powered assurance to detect anomalies and trigger automated remediation, and you can shrink your Mean Time to Resolution (MTTR) by 70%. That operational head room means organisations can run leaner teams yet maintain higher SLAs."

Collins adds that the right tools don't just reduce truck rolls — they raise the whole team's game.

"Tools that combine discovery, validation, and remote collaboration reduce the overhead of repeated site visits and manual data collection. Critically, efficiency gains come from choosing tools that integrate naturally into existing workflows and offer repeatable, standardised testing procedures. These tools not only reduce diagnostic time but enable more junior staff to perform tasks reliably, with the option for remote expert support when deeper analysis is needed."

Scaling without sinking the budget

With budgets under pressure, cost optimisation has become just as critical as technical performance. The challenge is showing that network upgrades aren't just necessary — they're strategic.

Burski cites cost savings from smarter architectures: "start by implementing SD-WAN to aggregate diverse connections (e.g., fibre, 5G, satellite) into a unified, policy-driven fabric. This reduces reliance on costly legacy infrastructure while enabling dynamic bandwidth allocation. Enterprises have cut operational costs by 30-40% by replacing fixed MPLS circuits with hybrid WANs that prioritise critical applications over affordable broadband during peak demand."

"Upgrading legacy devices isn't only about new features, it's a tidal wave of improved security and compliance. End-of-life platforms often miss critical firmware patches, leaving exploitable vulnerabilities," says Malik. "By migrating to modern infrastructure, complete with hardware root-of-trust, integrated threat feeds, and automated patching — you not only boost performance but also demonstrate due diligence to auditors."

According to Collins, ROI also shows up in what doesn't happen: "Demonstrating ROI is strongest when framed in terms of avoided costs (e.g. fewer outages, reduced escalations, improved resolution time) and support for business initiatives (e.g. hybrid working or secure remote access). It's not about overspending; it's about investing early to avoid reactive costs later."

Moreover, in a fast-moving tech landscape, buying a network solution is about more than specifications: it's about alignment and longevity.

"Successful integration isn't just technical — it's cultural," says Collins. "Tools that simplify onboarding, align with team workflows, and offer clear training paths foster faster adoption. Ongoing support should include open communication, regional responsiveness, and transparency especially as environments evolve."

"When evaluating new network technologies, enterprises should consider more than just technical specs; it must include ecosystem fit, roadmap alignment, vendor support maturity, and integration with existing tooling," advises Malik. "Don't just compare throughput or port counts, evaluate how well each solution dovetails with your existing operations frameworks, monitoring tools, and security policies. Engage vendors and independent partners early to set success criteria, such as session setup times, failover recovery metrics, or API-driven provisioning rates."

Hern sums it up: "Successful scaling requires businesses to take a strategic, forward-thinking approach to vendor selection and integration. It's not just about what technology is out there, but about choosing solutions that will support your business both now and in the future. The right decisions today will pay dividends tomorrow."

The future is flexible, automated — and already here

One truth unites all the expert perspectives: scaling today isn't just about adding more capacity. It's about building an enterprise network that is smart, secure, and responsive to constant change.

SD-WAN, orchestration, AI-driven insights, and cloud-native tools have become the new fundamentals. But it's the thoughtful integration of these elements — guided by strong partnerships and a sharp business lens — that will define the next era of enterprise connectivity ■

Hamzah Malik, CACI

Mark Burski, Digital Carbon

Avoiding common pitfalls in industrial Wi-Fi 7 deployments

Charlie McRae, Systems Engineer at IDS-INDATA

Wi-Fi 7 has the potential to significantly enhance connectivity in industrial environments, offering ultra-low latency and high-speed wireless performance that supports a range of applications, from smart factories to AI-enhanced automation.

In practice, industrial deployments often encounter difficulties when existing infrastructure isn't ready to meet the demands of Wi-Fi 7. Without careful planning and proper execution, businesses may experience downtime, suboptimal performance, and disappointing returns on their technology investments.

Below are three frequently encountered mistakes in deploying Wi-Fi 7 in industrial settings, along with practical ways to address them.

Mistake #1: Overlooking the importance of the wired infrastructure
Although Wi-Fi 7 offers significant performance improvements, its effectiveness depends heavily on the wired systems that support it. Many facilities still rely on legacy hardware, such as outdated switches or old Cat5 cabling, that can't handle the throughput required for newer wireless standards. When this foundation is neglected, even state-of-the-art access points cannot operate efficiently. The result is a bottleneck that reduces the impact of what should be a significant technological upgrade.
Impact: Essential operations can suffer delays or interruptions, ultimately reducing the effectiveness of digital transformation investments.

Mistake #2: Failing to meet power needs in harsh industrial settings
Wi-Fi 7 access points designed for industrial use often require more power than previous models, typically using Power over Ethernet Plus (PoE++) or IEEE 802.3bt standards. However, many facilities lack the necessary switching equipment, or their environments make it challenging to maintain consistent power.
Challenge: Supporting ruggedised access points with higher power demands in demanding environments requires thorough planning, rather than relying on off-the-shelf solutions.

Mistake #3: Ignoring RF challenges and 6GHz spectrum planning
Industrial environments often include heavy machinery, metal structures, and thick walls — all of which can interfere with wireless signals. With Wi-Fi 7 utilising wider 320 MHz channels and the 6 GHz spectrum, RF complexity increases considerably.
Result: Poor RF design can lead to reliability issues and service disruptions that negatively impact business-critical processes.

Laying the right foundation

A well-planned deployment starts with solid physical infrastructure. That means upgrading to shielded Cat6A cables and deploying multi-gigabit switches capable of handling Wi-Fi 7's throughput. Power delivery should be reviewed across the entire site, with PoE++-ready switchgear or industrial-grade injectors used where necessary.

To address environmental factors, Wi-Fi 7 access points should be rugged and rated for use in industrial spaces. Placement should follow detailed RF site surveys that consider interference, materials, and equipment layout.

Logical network design is also vital. Dividing networks to separate IT systems from operational technology (OT) traffic can preserve uptime and security. Regular monitoring, whether through managed services or in-house tools, helps ensure the network remains optimised over time.

Although WPA3 encryption is required for Wi-Fi 7 certification, relying on it alone to secure an industrial network can be a mistake. Many sites operate a mix of old and new devices, which can result in fallback configurations that compromise overall security.

New features, add functionality but also introduce new risks if not implemented securely. A lack of segmentation within the network can allow unauthorised access and lateral movement if a breach occurs. Without additional layers of protection, threats like rogue access points, device spoofing, or man-in-the-middle attacks remain possible, even in networks using WPA3.

Best practices

Security should be layered and proactive. Using certificate-based authentication (like EAP-TLS), enabling network access control (NAC), and applying Zero Trust principles can significantly enhance protection.

Separating IT and OT systems through microsegmentation can limit the impact of a breach. Wireless security reviews should be part of the same planning process as RF design to identify and eliminate vulnerabilities before deployment - a crucial step, and one not to be missed.

Wi-Fi 7 can revolutionise connectivity in industrial settings — but success depends on more than just installing new access points. ■

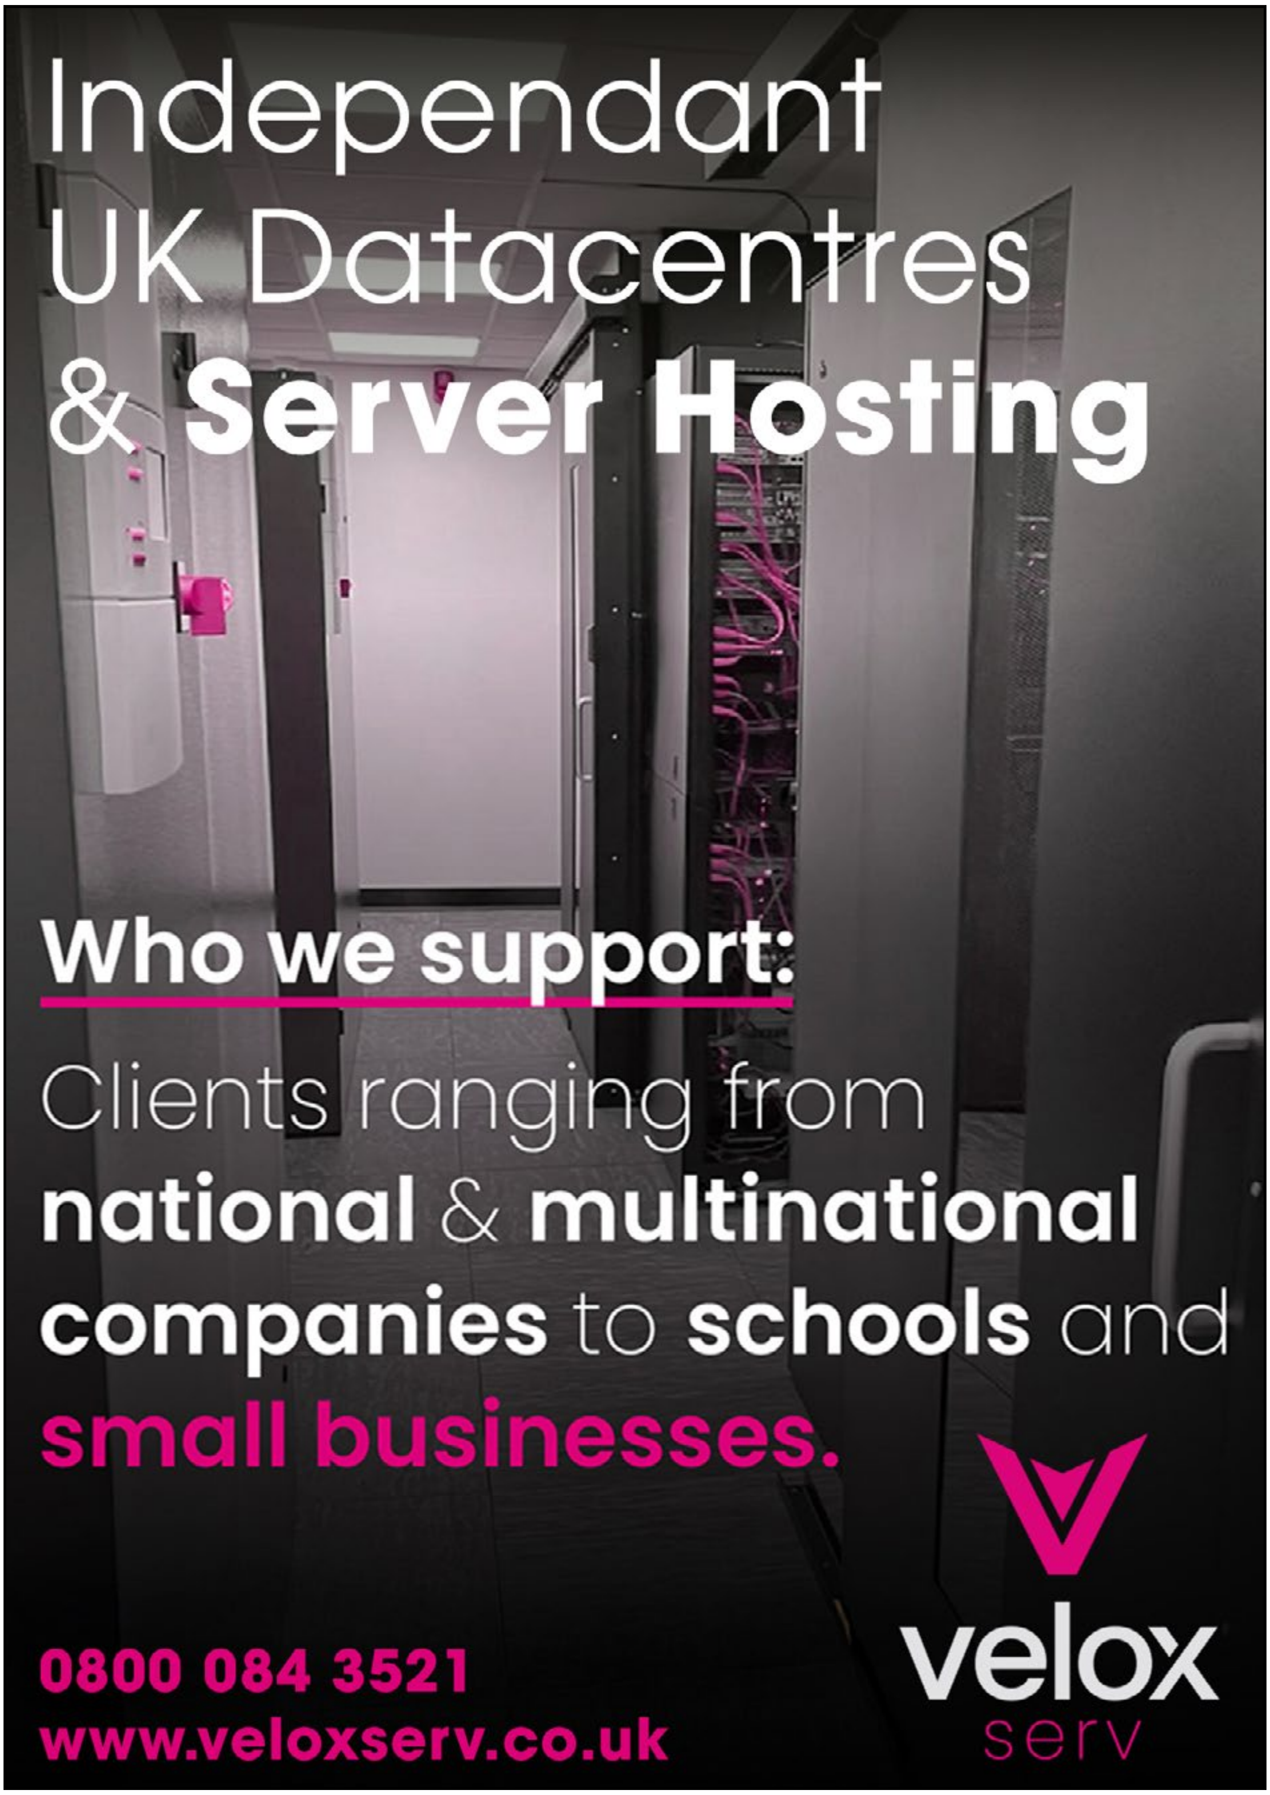

The new backbone of storage

Kubernetes has firmly established itself as the cornerstone of enterprise modernisation. Understanding its storage capabilities and challenges is essential…

Kubernetes is no longer an experimental technology; it's firmly entrenched as the backbone of enterprise modernisation.

"Kubernetes has reached mainstream adoption in the enterprise. It is now the standard platform for container orchestration, with widespread use in production across industries such as finance, healthcare, telecom, and manufacturing," says Ryan Kaw, VP of Global Sales at Catalogic.

"Kubernetes is the core of cloud-native modernisation, automating deployment, scaling, and managing containerised apps across complex, hybrid IT environments," agrees Divya Mohan, Principal Technology Advocate at SUSE and one of the maintainers of the Kubernetes project.

Moreover, enterprises are deploying Kubernetes not just in the cloud but also on-premises, as part of hybrid and multi-cloud strategies, supporting agile development, microservices architectures, and scalable infrastructure management.

Demystifying persistent storage

One of the most critical components for enterprise adoption is how Kubernetes handles persistent or non-volatile storage. In Kubernetes, persistent storage provides a way to store data beyond the lifecycle of a pod, ensuring that data is not lost when pods are recreated or rescheduled.

According to Mohan, "Kubernetes has fundamentally transformed how enterprises handle persistent storage by decoupling it from compute through the Persistent Volume subsystem, enabling stateful applications to run reliably and scale seamlessly."

Kaw explains that persistent storage in Kubernetes is managed via Persistent Volumes (PVs) and Persistent Volume Claims (PVCs), which abstract the underlying storage infrastructure. Kubernetes supports both static and dynamic provisioning using Storage Classes and Container Storage Interface (CSI) drivers, enabling integration with a variety of backends — from cloud-managed disks to on-prem enterprise systems.

"However, Kubernetes does not natively handle data protection tasks such as backup, restore, or disaster recovery. These responsibilities fall outside its core scope, requiring external tools, either commercial solutions or open-source options, to ensure workload durability and compliance," notes Kaw.

Compatibility, performance, and data protection

When designing Kubernetes storage solutions at scale, several key considerations come into play.

Oni Chakravartti, SVP, Head of Sales for Enterprise, Cloud Business Unit at Rakuten Symphony, lists the critical factors: "time to deploy, hyper-convergence, manageability from a single point across multiple regions, the ability to utilise multiple storage types, auto-scaling without disruption, high availability with three-way replication, and compatibility with existing and new hardware are all essential."

Since the goal of Kubernetes is to enable a portable, scalable, containerised architecture that supports changing application needs, in the enterprise, this must be supported by a scalable, secure, protected storage solution, that (ideally) can be delivered as software-defined, hardware-defined and cloud-defined.

"Organisations need to factor in reliability, scalability, performance, and security right from the start," confirms Mohan. "While Kubernetes offers the flexibility to automate and standardise storage with features like Dynamic Volume Provisioning and Storage Classes, it is imperative to complement them with robust backup and disaster recovery strategies to protect the data that is being stored. Given the ever-evolving nature of the threat and regulatory landscapes today, ensuring security and compliance is also non-negotiable, irrespective of the environment. Following best practices such as enforcing strict RBAC, encryption at rest and in transit, and regular key rotation will not only help organisations protect sensitive data but also meet compliance and regulatory requirements."

Mohan adds that while matching storage requirements with application needs helps with right-sizing storage resources, leveraging observability tools and automation to monitor storage utilisation, throughput, and latency can aid organisations in taking proactive and corrective measures.

However, integrating enterprise storage solutions with Kubernetes often encounters hurdles such as limited CSI driver support, architectural mismatches, operational overhead, and multi-tenancy issues. Mitigation strategies include using Kubernetes-certified CSI drivers with dynamic provisioning, integrating monitoring and alerting into cluster operations, and adopting Kubernetes-native storage projects for deeper control.

Kaw highlights the importance of vendor solutions that are purpose-built: "Bridging the gap between traditional storage paradigms and the dynamic nature of containers requires solutions that are designed for Kubernetes from the ground up, rather than relying on complex, risky open-source workarounds."

In contrast, Mohan believes that to bridge the gap between traditional storage paradigms and the dynamic, ephemeral nature of containers, open source is key.

"By investing in open source tooling that's purpose-built for such environments, not only can organisations leverage the benefits of collaborative innovation, but they can also save themselves from being locked to a cloud provider or vendor," asserts Mohan.

Cloud-managed storage: convenience with caveats

Managed storage services are popular for their ease of use, durability, and integration with cloud-native services. They support dynamic provisioning via CSI, providing seamless operations with minimal overhead.

"Using cloud-managed storage solutions with Kubernetes offers several advantages for enterprises, including minimal administrative overhead, automated scaling, built-in high availability, and seamless integration with cloud native monitoring and backup services," claims Mohan. "These platforms allow organisations to quickly provision persistent storage, benefit from the cloud provider's robust infrastructure, and focus on application development rather than storage management."

Chakravartti reports that "cloud solutions running on hyperscalers have inherent advantage of scale, however enterprises need data at the edge for use cases that need just-in-time inference and actions — this can only be done at the edge. The cost of moving data to cloud and bringing it back is high, and enterprises are seeing this cost continue to rise. For example, sending security files captured by camera to the cloud, processing the image and sending it back for actions does not deliver just-in-time needs. Also, the data movement is very expensive. Enterprises want to bring their cloud cost down, and the only way to do this is to send and store data that is needed on cloud, and inference and actionable data to be processed at the edge."

And cost isn't the only issue; enterprises may face limitations in control and customisation, as managed services restrict infrastructure choices and integrations to what the provider supports.

"Organisations risk vendor lock-in, making future migrations or hybrid strategies more complex. Additionally, meeting strict compliance or data residency requirements can be challenging, as complete visibility and control over the underlying storage infrastructure are limited," warns Mohan.

According to Kaw, "to address these drawbacks, many organisations adopt cloud-agnostic backup solutions or extend open-source tools to manage data across environments, ensuring consistent backup, recovery, and migration workflows."

Ready for the next step?

Kubernetes is fundamentally reshaping how enterprises think about storage — moving from static, siloed solutions to flexible, scalable, and cloud-native architectures.

And, looking ahead, the enterprise market is expected to accelerate its migration to Kubernetes.

"We expect enterprise to migrate 70% of apps onto Kubernetes platform in the next 2-3 years. Supporting this trend, independent research from IDC forecasts that by 2027, more than 75% of all AI deployments will leverage container technology, and by 2028, 80% of custom software at the edge will run in containers — up from less than 50% today," says Chakravartti.

As enterprises continue their digital transformation journeys, understanding and harnessing Kubernetes' storage capabilities will be pivotal. Embracing open standards, leveraging hybrid architectures, and partnering with experienced vendors will help organisations unlock the full potential of containerised enterprise applications. ■

The AI gold rush is on — why storage MSPs hold the key to infrastructure success

Paul Speciale, Chief Marketing Officer, Scality

In the 2020s, we've crossed a rubicon in the AI revolution. What were once long-discussed concepts about the possibilities and pitfalls of AI have exploded into reality. With generative tools leading breakthroughs in content creation, data analysis, and coding, the market reflects this momentum - AI is set to soar from $93 billion in 2020 to $826 billion by 2030. Moreover, key analysts like Gartner now predict a massive growth in enterprise demand for consumption-based as-a-service based offerings, further amplifying the market opportunity.

However, with great potential comes great responsibility. The pressure is now on for managed service providers. As stewards of digital infrastructure, MSPs must go beyond the mere baseline of provisioning resources. They need to entirely reimagine their role in helping clients harness AI effectively. But this key challenge is also a chance - and nowhere is this more true than in today's dynamic AI landscape. If MSPs play this right, they can turn this moment into a golden opportunity, expand their service offerings and capitalise on the growing needs around AI.

> **"As stewards of digital infrastructure, MSPs must go beyond the mere baseline of provisioning resources. They need to entirely reimagine their role in helping clients harness AI effectively."**

Legacy storage is tired

Traditional infrastructure approaches - especially legacy storage systems - are now totally inadequate for today's demands. They are not designed to handle the unpredictable, dynamic, high-throughput demands of AI workloads. To remain relevant and create long-term value, MSPs should therefore opt for super-agile, software-defined, multidimensional scaling solutions. This approach empowers them to scale infrastructure independently across multiple axes, providing the flexibility that AI workloads demand.

As AI reshapes industry expectations, clients want partners who understand the strategic value of AI and can architect infrastructure that accelerates innovation.

To step into this expanded role, MSPs must evolve from service providers into strategic AI advisors. This means investing in a highly flexible, scalable, intelligent infrastructure that aligns with business outcomes - whether enabling real-time analytics, streamlining data governance, or scaling AI model training environments. By adopting infrastructure models that prioritise flexibility and performance, MSPs can directly support their clients' AI-driven transformations and secure their own growth in the process.

Delivering hyperscaler agility in the private cloud

Many organisations seek the elasticity of public clouds but require the data control and compliance guarantees of private environments. MSPs can bridge this gap by deploying private cloud platforms that emulate the agility of hyperscalers. With automated scaling, user-friendly interfaces, and rapid provisioning, MSPs can meet client expectations while ensuring data remains secure and localised.

Packaging scale with compliance as a unified offering

AI workloads are inherently unpredictable, with spikes in data usage and performance needs. Through multidimensional scaling, MSPs can fine-tune their infrastructure - ramping up resources only where needed. This not only prevents overprovisioning but also ensures that sensitive data remains compliant with industry regulations, such as GDPR or HIPAA.

Going local: optimising infrastructure for data sovereignty

As data privacy regulations tighten globally, localised infrastructure is becoming a necessity. By deploying regional cloud offerings, MSPs can help clients meet national data residency requirements while delivering low-latency performance. This localised approach is not just a compliance measure - it's a strategic advantage.

Supporting consumption-based and multi-tenant models

The shift to AI is accelerating the demand for flexible billing models. MSPs should offer consumption-based pricing and multi-tenant architecture to accommodate the bursty, iterative nature of AI development. This ensures clients can scale up or down based on actual usage, improving satisfaction while maintaining cost transparency.

MSPs that proactively support AI workloads with tailored infrastructure are positioned to unlock significant new revenue streams. AI is resource-intensive, and clients are seeking partners who can meet these requirements with resilient, high-performance solutions.

Capturing AI infrastructure spend

The demand for compute and storage is rising in tandem with AI adoption. MSPs that offer AI-optimised SLAs, scalable capacity, and high-throughput processing will stand out as preferred partners in this growing market.

With multidimensional scaling, MSPs can deliver infrastructure that precisely matches workload demands. Whether scaling up storage for massive datasets or reducing latency for inference workloads, this tailored approach boosts efficiency and protects margins - turning infrastructure from a cost centre into a strategic asset.

Why now is the time to embrace modern storage offerings

The age of AI demands a new playbook for infrastructure - and MSPs have a pivotal role to play. By embracing software-defined, multidimensional storage scaling, MSPs can provide the flexibility, performance, and compliance clients need to succeed in a data-driven world.

This transformation isn't just about technology — for MSPs, it's a golden growth strategy. Those who act now will be better positioned to serve AI-driven businesses, open up new revenue channels, and evolve into indispensable partners in digital innovation. MSPs that welcome the challenges of AI today will lead the market tomorrow. ■

Offshore equipment specialist receives future-proofed communications

Whether it's fossil fuels or renewables, the global offshore energy sector is growing. Oil demand is forecast to be 3.2 million barrels per day higher in 2030 than in 2023. Meanwhile, global offshore wind is set to increase from 80GW today to an estimated 212GW by 2030.

Offshore sites involve many stakeholders and face unique operational challenges and regulatory demands. Because of this, the offshore energy sector already has several distinct communication needs. So, as demand increases over the coming years, so too will the complexity of communication between the many stakeholders in the sector.

It is against this backdrop that a UK-based specialist offshore equipment provider needed to streamline communication between hundreds of employees while also enhancing its ability to connect with clients globally. Failing to do so would get in the way of its operations and potential growth. With an eye on future scalability, CallTower created a Unified Communications as a Service (UCaaS) and PSTN solution to help.

Supplying specialised equipment around the world

The UK-based equipment provider delivers specialised offshore equipment to locations around the world, along with asset inspection and maintenance services. Each location and market has unique needs and regulations, and although a mid-market company, it has a presence across Europe, Africa, and Australasia, for effective project delivery. However, to ensure it can effectively communicate with stakeholders across these markets, the equipment provider needed to enhance its telephony capabilities so its team could seamlessly connect with clients and stakeholders wherever they were.

"Offshore energy sites are almost always joint ventures between large organisations with significant amounts of capital invested. This makes stakeholder alignment a mission-critical process involving many players. Organisations like our client's can't risk having poor connectivity in an environment like this," said CallTower CRO William Rubio.

Global operations come with global challenges. Managing teams and operations in eight countries, the equipment provider required a robust communication and collaboration solution to support its own geographically dispersed workforce, too. It needed to streamline its internal communications and collaboration between hundreds of employees across multiple time zones.

On top of this, aware of the growth of the sector, the equipment provider wanted a solution that was scalable. It was important that any solution could meet future growth in demand and changing business needs.

"The way the world sources energy is changing, and companies in the energy sector need communications systems that can keep up," said Rubio.

Combining technologies in a single communication platform

To meet the equipment provider's external and internal communication needs, CallTower helped create and deploy a solution combining Microsoft Teams and high-quality PSTN calling integration. The Teams solution was provisioned using E5 licenses.

This combination of technologies is housed in a single, unified communications platform, designed for operational efficiency, ease of use and future scalability. It also features the ability to add mobile eSIMs to allow team members to call colleagues, clients or partners when they're away from their desk.

The Microsoft Teams deployment means the equipment provider's workforce can now collaborate through voice, video, and messaging from within a single, familiar tool. Teams in Europe, Africa, and Australasia can communicate more easily, fostering improved connectivity across time zones. Thanks to the PSTN integration, calling global stakeholders and clients is seamless and effortless as well.

By combining Teams and PSTN calling into a native-like application, there's also no need for multiple communications tools, reducing the time and risk of swapping between, and moving data between, applications. The unified platform also simplifies the equipment provider's setup, making the platform easier to for IT teams to use and manage.

Mobile eSIM capabilities equip the provider's communications infrastructure to support evolving business needs. As new team members join, the platform can be easily scaled, and it supports bring-your-own-device (BYOD) for those who need to work on location or offshore.

"Though the equipment provider's communication needs became more complex, its communication systems were simplified. Its employees can better collaborate and manage stakeholder relations, which we know is of growing importance in the offshore energy sector," said Rubio. "Global energy demand is set to grow 11-18% between now and 2050. Our solution allows the equipment provider to grow and more easily adapt to new scenarios and market requirements. Organisations that can't adapt are at risk of being left behind." ■

Thames Water transforms asset management with Getac

Thames Water is in the midst of a multi-year organisational transformation to optimise its asset management and customer service capabilities through the adoption of innovative digital technology. A key part of this is the migration of its field-based workforce to a flexible new Android-based ecosystem.

An Android migration

As the UK's largest private utility company, Thames Water is responsible for extensive water management infrastructure across London and the Thames Valley in South-East England. Per day, the company supplies 2.5 billion litres of drinking water and treats 4.6 billion litres of wastewater, for 15.5 million people (about a quarter of the UK population).

Thames Water is currently undergoing a major, multi-year organisational transformation, using the latest digital technology to future proof its operations and bolster both its asset management and customer service capabilities.

A key part of this transformation is the migration of its field-based workforce over to a flexible new Android-based ecosystem, giving engineers fast and reliable access to critical information at the touch of a button - even when working in remote or challenging environments. Doing so enables them to operate more efficiently, solve complex maintenance issues faster, and deliver a higher number of first-time fixes to customers.

"With advantages in cost efficiency, energy efficiency, variety and customisation, Android OS allows us to tailor IT solutions to our exact needs," says Ade Ayajo, Senior Project Manager, Thames Water. "After making the decision to migrate over to Android, we knew we needed to equip our field engineers with new digital devices that could maximise all the benefits Android has to offer when out maintaining field assets and during on-site customer visits."

Owing to the multi-year nature of the project, Thames Water also wanted to work with a solutions provider that could deliver long-term support in key technical areas including the devices themselves, OS versions and security patches.

Driving rapid transformation

Thames Water assessed numerous solutions and vendors before settling on a fully rugged solution built around Getac's ZX10 fully rugged Android tablet, which combines powerful performance, reliability and long-term security support in a slim and versatile profile.

As digitisation continues to drive rapid transformation across the utilities sector and companies replace outdated paper-based processes with faster digital alternatives, the need for fully connected field teams has become more important than ever. Field engineers and technicians require multipurpose devices that they can use to accomplish a diverse range of tasks simultaneously, from capturing data and monitoring critical assets, to diagnosing issues in customers' homes and remotely collaborating with colleagues or subject matter experts. Without the right devices this becomes incredibly difficult to achieve, leading to service delays, unplanned downtime, and lower levels of customer service.

On top of this, the devices need to excel in the challenging environments that utilities field engineers often work in. This includes regular use in changeable temperatures and weather conditions (including rain and snow), working above and below ground, as well as withstanding accidental knocks, bumps and drops – all of which are part of everyday field work. If devices aren't up to the task, damage and failure rates quickly start to rise, significantly impacting the total cost of ownership (TCO) over time.

The Getac ZX10 is designed from the ground up to provide exceptional reliability in the difficult operating environments that utilities field engineers find themselves in every day, resulting in incredibly low TCO for each device. Another key reason why Thames Water chose the ZX10 is because it is an Android Enterprise Recommended device, which means it is guaranteed to receive a minimum five years of security patch support from its release date, offering long-term peace of mind.

The Android support strategy proposed by Getac not only provided the required reassurances of being able to support the devices with 90-day security updates until 2030, but it also means Thames Water can take advantage of new Android features in a timely manner.

Increased satisfaction

With the new rugged solution now in the hands of its field engineers, Thames Water is already realising the benefits of its new Android ecosystem, including less downtime, faster resolution of issues, and higher levels of customer satisfaction.

"We've had fantastic feedback from across our field-based workforce," says Ade. "The Getac's devices are a huge asset to the team, offering new levels of performance and reliability in a wide range of scenarios. Whether conducting maintenance at our water treatment facilities, or working on site with customers, our engineers know they have both the hardware and software they need to get the job done quickly and efficiently."

In addition to providing the devices themselves – along with peripherals including detachable keyboards and hand straps – Getac is also working closely with Thames Water to deliver ongoing support in the key areas it requires, including regular Android security patches, device maintenance, OS updates and Salesforce integration.

"This Android migration is a long-term project for Thames Water, so we knew we needed a partner that would be with us every step of the way," adds Ade. "In addition to providing class-leading hardware and peripherals, Getac's technology experts have also integrated seamlessly into our own team. Should problems arise, we know we can count on them to find fast resolutions and keep our field force operating at an optimal level, giving us total peace of mind." ■

As AI gains momentum, we'll all be talking more about "neural edge"

Peter Wilcock, VP, Latos Data Centres

At Nvidia's GTC developer conference in March, CEO Jensen Huang heralded a trillion-dollar boom in data centres, to underpin ever-more sophisticated AIs. He declared that "the data centre is no longer a warehouse for computing, it is the engine of AI."

This transformation is already underway. Since 2021, sales of conventional CPUs have slumped by 80%. By contrast, demand for the GPUs needed for AI is accelerating, and expected to grow by almost 30% every year.

As AI moves from experimental to essential, the traditional architecture of our data centres — concentrated, centralised, and physically distant — is no longer fit for purpose.

For the UK to reap the full benefits of AI, we need to rethink the country's data centre map.

The real-time AI revolution

Up to now, hyperscale data centres are designed to retrieve data – files, software processes, and so on. The "engines of AI" Jensen Huang envisions are designed from the ground up to deliver the processing power to interpret, learn from, and respond to new information on the fly – and at massive scale.

This revolution in real-time AI is set to transform how we live and work. But sheer computing heft isn't the whole answer. Where that computing is located is going to be increasingly important.

Take the metaverse, gaining ground in home entertainment and industrial training environments alike. Latency is a make-or-break consideration: delays longer than 20 milliseconds or so can cause disorientation and motion sickness for users. That is the maximum delay tolerable for AI-powered content, facial tracking, or environment rendering.

Transportation is another example. Autonomous vehicles depend on decisions made in microseconds to ensure passenger and pedestrian safety. A car that needs to ping a distant data centre won't respond quickly enough to a sudden hazard. The slightest delay could be a matter of life and death.

In predictive healthcare, AI systems analyse patient data to provide early warnings and enable tailored interventions. There have been many encouraging advances in the use of AI to identify the early stages of cancer, Alzheimer's disease, COPD, and other conditions. Not only must predictive healthcare AIs be clinically accurate, they must also deal with a wide variety of data – images, videos, and written records – with the utmost security.

In other words, extent to which an AI can be valuable, responsive, and safe, will be increasingly measured in kilometres – from the end-user to the data centre.

Delivering on its promise depends on complementing existing hyperscale infrastructure with new facilities, designed for AI, located much closer to end users.

A new generation of edge

A new generation of smaller scale 'neural edge' data centres are coming onto the market, focused on meeting precisely this need.

What sets these facilities apart is their energy-dense design – needed to support the most demanding AI training and inference workloads.

But performance specs are only part of the story. Neural edge data centres can be built quickly and integrated unobtrusively into urban environments. Modular building techniques and sustainable power and cooling mean facilities can start small and grow in line with demand, without compromising their overall performance.

By placing compute power close to end users, in a tech cluster, a manufacturing hub, or even a residential area, neural edge facilities can dramatically reduce AI latency while improving overall resilience. And they are likely to prove more energy efficient in the long run.

The UK's AI moment

The UK's ambition to lead in AI is not just about innovation policy. It's about building the physical foundations for an AI-driven economy – expected to add almost £50 billion to the economy every year. The government's promotion of AI Growth Zones — regions earmarked for accelerated investment in AI infrastructure — reflects an understanding that AI must be supported regionally, not just centrally.

If that adoption is to be inclusive, scalable, and sustainable, it must be backed by infrastructure that ensures real-time AI is available where it's needed.

It's easy to get swept up in the promise of AI — the breakthroughs in healthcare, the efficiencies in transport, the transformations in everyday life. But without the right infrastructure, those breakthroughs risk stalling. Neural edge data centres offer a practical, scalable way to deliver real-time AI where businesses, public services, and homes need it.

If AI is the brain of the digital economy, the neural edge is its nervous system. Fast, distributed, and always on, neural edge data centres will enable the next decade of UK innovation — not in the cloud, but from the ground up. ■

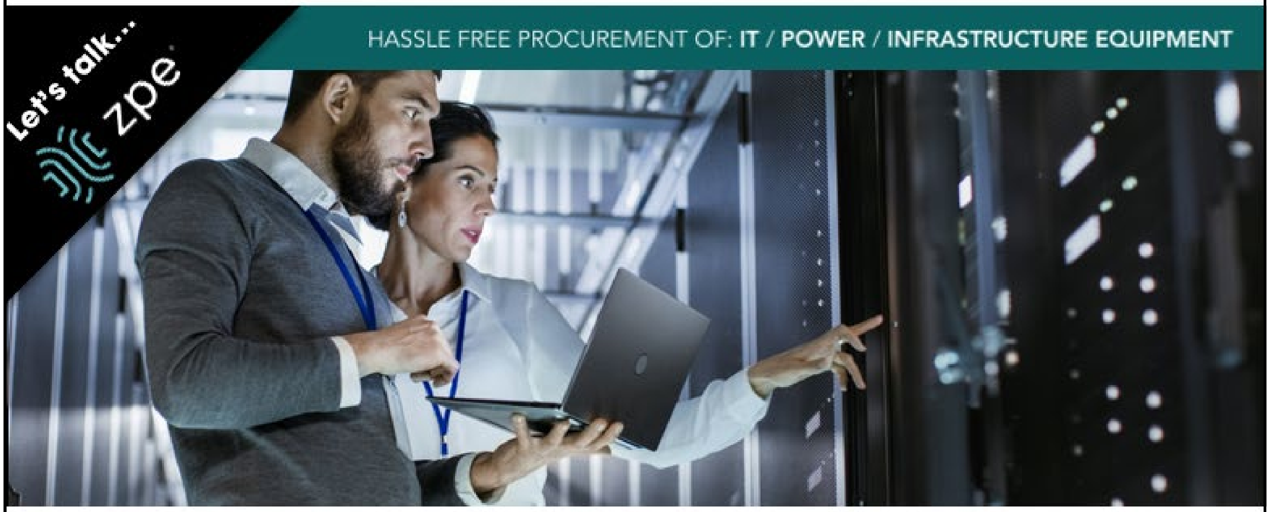

Network security solutions: selection and deployment

Matthew Terry, Senior Product Manager, WatchGuard Technologies

Cyber adversaries grow more agile and sophisticated every day. This means that the level of security that an enterprise needs, has to go far beyond simply installing traditional firewalls and antivirus software. All UK enterprises these days are dealing with advanced malware and ransomware, as well as persistent, targeted attacks.

With that in mind, some top tips for choosing a hybrid network security solution, that fully protects remote and office users, include:

First fully assess your infrastructure: Before selecting any solution, it's essential to thoroughly evaluate current infrastructure. It is important to look at factors such as the number of endpoints, the degree of cloud adoption, and what existing security tools are. Organisations with distributed offices or those handling more sensitive data need advanced AI-driven threat detection and real-time response capabilities.

Embrace AI-driven, Zero-Day protection: In the era of polymorphic malware and zero-day exploits, traditional signature-based antivirus is no longer enough. Look for a solution that leverages AI-powered threat detection and behaviour-based analysis. To be as secure as possible, potential threats need to be identified and neutralised in real time. This is the only way to ensure your network stays one step ahead of ever-evolving risks.

Look to unified threat management: Fragmented security systems all too often create gaps and inefficiencies for criminals to exploit. By adopting a unified security platform that marries next-generation firewall capabilities, Intrusion Prevention Systems (IPS), secure web gateways and centralised management, UK enterprises get to view and manage their security posture through a single-pane-of-glass. This allows for faster threat correlation and accelerates incident response across the network.

Use cloud-enabled scalability and real-time threat intelligence: Remote work and cloud-based infrastructures are commonplace in most organisations which means scalability and agility are essential. Cloud-based solutions provide the flexibility and rapid threat intelligence updates that are needed to defend against emerging vulnerabilities. Having a flow of real-time intelligence ensures that your defences evolve at pace with the threat landscape.

Ensure compliance with regulatory standards: With stringent compliance requirements such as GDPR, Cyber Essentials, and the NIS2 directive, your network security must meet these standards. Find a solution that not only offers robust protection but also integrates regulatory considerations. This will help you avoid costly penalties and safeguard your reputation.

Foster a culture of security awareness: Human error is still one of the most prevalent causes of cybersecurity breaches and no technology can overcome human error alone. Regular, targeted security awareness training is critical to harness the technology and guard against social engineering tactics and phishing scams.

Performance without compromise: Security measures available in a hybrid network security solution shouldn't negatively impact your organisation's productivity. Ensure your solution has minimal impact and provides seamless performance so your workforce is productive even in the most secure environments.

Conduct pilot deployments: Before a full rollout, pilot testing the solution in a controlled segment of your network can provide valuable insights. This allows you to evaluate performance and ease of integration so you can fine-tune configuration to reduce false positives. It helps to include key stakeholders early to ensure that when deployed enterprise-wide, the solution actually meets their needs and your operational needs.

Robust cybersecurity hinges on a proactive, unified defence strategy that goes beyond isolated tools like anti-virus.

Cybersecurity solutions designed for distributed enterprises: The hybrid model workspace concept is at the forefront of business operations in today's rapidly changing work environment. As companies transition to flexible work arrangements and remote collaboration, ensuring secure access to corporate networks, devices, and data has never been more critical.

The shift to hybrid work models has increased the opportunities available to cybercriminals. Both employees and work devices between on-site and remote locations may present a security risk.

Hybrid work models empower employees to balance remote and in-office working arrangements, choosing where and often when they work. Such a significant change requires businesses to protect their hybrid workplace from cyberattacks with a modern security stack.

Unified security is the key to comprehensive protection. It acknowledges that no single security product or solution is infallible. Instead, it combines various security measures to safeguard your client environments, devices, and users. By creating multiple layers of defence, we reduce the gaps and vulnerabilities between each layer, making it significantly more challenging for cybercriminals to exploit weaknesses. ■

PRODUCTS

Bitdefender GravityZone is an enterprise-grade endpoint security platform that provides comprehensive protection against malware, ransomware, zero-day exploits, and advanced cyber threats. Built on a scalable, cloud-based architecture, GravityZone offers centralised management, real-time threat detection, and automated response capabilities across diverse endpoints, including desktops, servers, and virtual environments.

The platform employs multiple layers of security, including machine learning-based threat detection, behavioural analysis, and signature-based antivirus. Its advanced threat defense features include Exploit Prevention, Anti-Phishing, and Ransomware Remediation, which proactively block malicious activities before damage occurs. GravityZone's advanced machine learning models analyse threat behaviour patterns to identify zero-day threats with high accuracy.

GravityZone also integrates Endpoint Risk Management, providing vulnerability assessment, patch management, and application control to reduce attack surfaces. Its Sandbox Analyser isolates suspicious files for in-depth inspection, ensuring zero false positives. The platform offers flexible deployment options — on-premises, cloud, or hybrid — and supports Windows, macOS, Linux, and virtualised environments.

Management console dashboards deliver detailed security analytics, threat reports, and policy enforcement, enabling security teams to respond swiftly to incidents. Certified compliant with standards such as GDPR, HIPAA, and PCI DSS, Bitdefender GravityZone ensures robust, scalable, and compliant endpoint security.

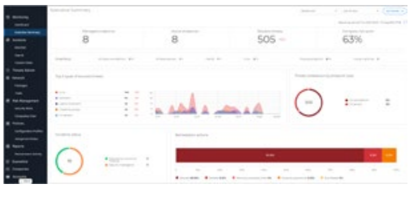

Proofpoint Security & Compliance is an advanced enterprise security platform designed to protect organisations from sophisticated email threats, data leaks, and regulatory violations. Leveraging cloud-native architecture, it provides real-time threat detection and data loss prevention (DLP) across email, social media, and cloud applications.

The platform utilises machine learning algorithms, behavioural analytics, and threat intelligence to identify and block advanced phishing, malware, and impersonation attacks. Its email security features include URL rewriting, attachment sandboxing, and adaptive scanning, ensuring zero-hour threat prevention. Proofpoint's Targeted Attack Protection (TAP) uses sandboxing and threat intelligence feeds to analyse suspicious content in real-time.

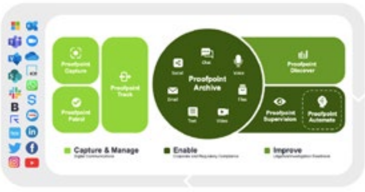

For compliance, Proofpoint offers comprehensive DLP capabilities, monitoring content across email and cloud services to prevent accidental leaks of sensitive data such as PII, PHI, or intellectual property. It supports policy enforcement with customisable rules, encryption, and audit logging, ensuring regulatory adherence to standards like GDPR, HIPAA, and PCI DSS.

The platform integrates seamlessly with Microsoft 365, Google Workspace, and other cloud platforms, providing centralised management through a unified console. Its threat intelligence cloud constantly updates protections, enabling organisations to stay ahead of evolving cyber threats while maintaining compliance and data security.

Check Point Infinity is a unified cybersecurity architecture designed to provide comprehensive threat prevention across networks, cloud, endpoints, and mobile environments. Built on a scalable platform, Infinity integrates advanced security solutions like Threat Prevention, SandBlast Zero-Day Protection, CloudGuard, and Harmony Mobile into a single, cohesive system.

The platform leverages next-generation firewalls (NGFW) with Stateful Inspection, application control, and intrusion prevention system (IPS) capabilities, utilising ThreatCloud AI and threat intelligence for real-time threat detection. Infinity supports multi-vector security with full SSL inspection, ensuring encrypted traffic is analysed without compromising performance.

For cloud security, Check Point Infinity integrates with major cloud providers, offering posture management, workload protection, and identity-aware security controls. Its endpoint security includes sandboxing, anti-malware, and device control, protected via centralised management via the Infinity Portal.

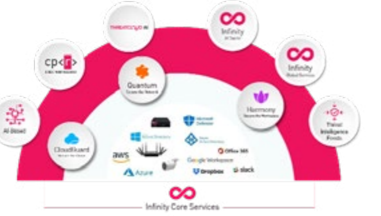

The architecture emphasises scalability and automation, supporting multi-domain management and threat response orchestration. It complies with industry standards such as GDPR, HIPAA, and PCI DSS, ensuring regulatory adherence. The platform's threat intelligence, combined with AI-driven analytics, delivers proactive security, reducing attack surfaces across hybrid environments.

The **Huntress Managed Cybersecurity Platform** is a proactive threat detection and response solution designed for Managed Service Providers (MSPs) and security teams. Built on a lightweight agent architecture, Huntress continuously monitors endpoints, servers, and cloud workloads for malicious activity, enabling rapid detection and remediation of advanced persistent threats (APTs), ransomware, and other sophisticated attacks.

The platform leverages behavioural analysis, threat hunting, and threat intelligence integration to identify hidden threats that bypass traditional security controls. Its detection engine analyses system behaviours, file activity, and network connections in real-time, flagging anomalies indicative of malicious activity. Huntress's threat hunting capabilities empower security teams to proactively seek out threats based on emerging attack techniques.

Huntress provides automated remediation workflows, including threat quarantine and removal, reducing incident response times. Its centralised dashboard offers detailed insights, attack timelines, and comprehensive reporting, facilitating compliance and audit requirements. The platform supports integrations with SIEMs and ticketing systems for streamlined workflows.

Designed for MSPs and security teams, Huntress's lightweight agent ensures minimal performance impact while delivering enterprise-grade threat detection. Its cloud-based management console simplifies deployment, monitoring, and threat analysis for organisations seeking proactive cybersecurity defense.

"Please meet...

Dominic Mensah, Director, Strategic Accounts, Lakeside Software

Who was your hero when you were growing up?

I look up to people who inspire curiosity, creativity, and kindness; for example, Marie Curie, who broke barriers in science (becoming the first woman to win a Nobel prize), and Nelson Mandela, who spread forgiveness and reconciliation.

What was your big career break?

My first career break was landing a role as a support engineer for a theatre booking platform. It taught me a lot about patience, active listening, troubleshooting computers, and understanding the challenges users face with their devices and applications while trying to stay productive. This experience cemented my love of technology, and opened my eyes to the complexity of IT operations and their impact on customer experience, leading me to where I am today.

What did you want to be when you were growing up?

I wanted to be an explorer - setting off into uncharted lands, discovering hidden treasures, and learning about new cultures. Now, in my current job, I channel that same spirit of exploration as I delve into the impact of technology and network integrations on workplace dynamics. My focus on Digital Employee Experience (DEX) allows me to uncover how these integrations enhance employee satisfaction, which is crucial for maintaining engaged and productive teams.

What's the best piece of advice you've been given?

Be more interested than interesting. This means focusing on curiosity, listening deeply, and genuinely engaging with others rather than just trying to impress them. It fosters stronger relationships and continuous learning. It's this focus on being a good colleague and maximising my knowledge base that has resulted in job satisfaction each and every day. I am very lucky to work somewhere where the be more interested than interesting methodology works perfectly.

If you had to work in a different industry, which would you choose?

Top of my list is music production. The idea of using AI to create, mix, and produce music fascinates me. Music is both an art and a science, blending creativity with technology. I'd love to experiment with AI-generated compositions and collaborate with artists, real or virtual, to enhance their sound. But perhaps the reason behind my interest is more the use of AI in orchestrating a future where music is personalised. AI has such great potential to change so many different industries, meaning music isn't the only one that will benefit from that technology moving forward.

What would you do with £1 million?

Invest in an AI company, as artificial intelligence is rapidly transforming the digital landscape. The next big thing that AI is going to be invaluable for is self-healing technologies, for example, PCs with smart agents that can fix issues and ensure they run smarter, faster, and more efficiently. In fact, we've been working with Intel and the firm has already unveiled neural processing units that combine with our endpoint monitoring platform to optimise performance, and IT management, and enable predictive analytics and maintenance.

Where would you live if money was no object?

I'd fully immerse myself in experiencing different cultures by travelling and living in diverse places around the world. The people, landscapes, cultures and traditions fascinate me. I imagine how fun it would be to live in far flung places like Thailand, Vietnam, the Galapagos and even the Antarctic. I live in the UK and would love to visit America more often to see my colleagues over there, as Lakeside Software has offices in Boston and Atlanta.

What's the greatest technological advancement in your lifetime?

The rise of Artificial Intelligence. AI has transformed countless industries, from healthcare and finance to entertainment and space exploration. The ability of AI to process vast amounts of data, automate complex tasks, and even engage in human-like conversation has reshaped how people live and work. AI adoption in companies has not only automated business processes but also acts as a digital colleague for bettering employee experiences. Breakthroughs in deep learning, natural language processing, and robotics continue to push the boundaries of what's possible – in proactive and predictive IT, digital employee experience, self-healing AI PCs, and much more. ■ "