JUNE 2021
networkingplus.co.uk

NETWORKING+

Cybersecurity is not a one-stop-shop

Security should be reinforced

Giacom and Webroot explain, p6

Separating the wheat from the chaff

Alert noise - a big IT Issue

Martin Hodgson, Paessler, p7

The ‘hybrid cloud visibility gap’

What it is and why it matters

Adrian Rowley, Gigamon, p13

Internet outage caused by software bug

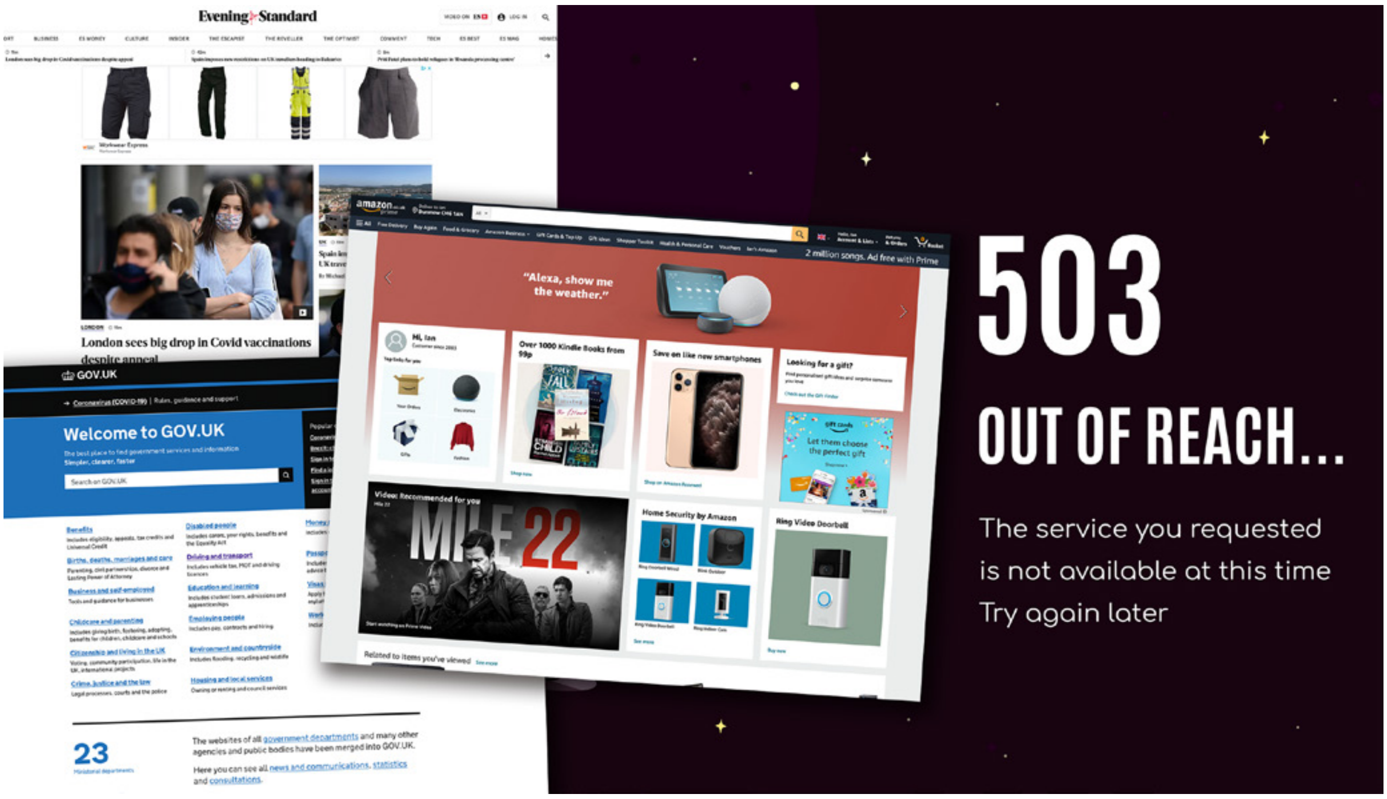
Evening Standard
London sees big drop in Covid vaccinations despite appeal
GOV.UK
Welcome to GOV.UK
The best place to find government services and information
Simpler, clearer, faster
23
Ministerial departments
The websites of all government departments and many other agencies and public bodies have been merged into GOV.UK.
Here you can see all news and communications, statistics and consultations.
amazon
"Alexa, show me the weather."
MILE 22
503
OUT OF REACH...
The service you requested is not available at this time
Try again later

The global internet blackout that hit many high-profile websites earlier this month, including the UK government’s, has been blamed on a software bug.

Fastly, the US cloud-computing company responsible for the issues June 8, apologised and said the bug had been triggered when one of its customers had changed its settings.

It is used by many major companies to help speed up loading times by storing versions of their websites in local servers.

The outage, which lasted about an hour, hit some popular websites such as Amazon, the Guardian, Evening Standard the New York Times and Reddit.

The UK government’s website going down added to the frustration of thousands of young people who were already struggling to book their Covid-19 vaccines after the 25-29 age group became eligible.

A day later, the Fastly’s senior vice president of engineering and infrastructure, Nick Rockwell, wrote a blog, in which he explained what happened and apologised for the incident.

“We experienced a global outage due to an undiscovered software bug that surfaced on 8 June when it was triggered by a valid customer configuration change,” he wrote. “We detected the disruption within one minute, then identified and isolated the cause, and disabled the configuration. Within 49 minutes, 95% of our network was operating as normal. This outage was broad and severe, and we’re truly sorry for the impact to our customers and everyone who relies on them.”

However, a customer quite legitimately changing its settings had exposed a bug in a software update issued to customers in mid-May, causing “85% of our network to return errors”, it said.

Avinash Prasad, head of managed security services at Tata Communications, said “the paradigm shift has happened from conventional viewing to predominantly internet-based viewing, from text to video and from recorded to live and interactive”, which is increasingly putting pressure on network infrastructures of service providers, media houses and online businesses among others.

“To cater to this additional demand, the technology and service provider industry is innovating to deliver high-quality viewing experiences, while optimising bandwidth usage, and creating new revenue streams,” he said. “A reliable and experienced Content Delivery Network (CDN) service enables media players to optimise infrastructure distribution and costs, allowing for maximum uptime and channel scalability, ensuring a seamless end-user experience.

Mark Hendry, director of data protection and cyber security at global legal business, DWF added: “The intention of CDNs is to route (or distribute) internet traffic and services through ‘nodes’ in order to balance the load of traffic, prevent bottlenecks and result in high availability and faster content delivery. Requests for content are directed by an algorithm, for instance the algorithm might direct the traffic so that it routes through the most available or highest performing node, or so that the traffic takes the fastest network route to the requestor.”

The outage has also raised questions about relying on a handful of companies to run the vast infrastructure that underpins the internet. ■

IN DEPTH:
SD-WAN products focus
P11-12

Ransomware attacks down as criminals focus on more lucrative targets

Ransomware attacks dropped by 50% in Q1 2021 as threat actors moved from using mass spread campaigns to focusing on fewer, larger targets with unique samples.

That is according to the *McAfee Threats Report: June 2021,* in which researchers noted that the traditional approach of using one form of ransomware to infect and extort payments from many victims is becoming less prominent. The reason for this, the researchers said, is because the targeted systems can recognise and block such attempts over time.

Instead, a trend is developing toward fewer, customised Ransomware-as-a Service (RaaS) campaigns aimed at blue chip organisations.

The research identified that the number of prominent ransomware family types declined from 19 in January 2021 to nine in March 2021.

Furthermore, the report found that the most detected ransomware group in the first quarter of this year was REvil, followed by RansomeXX, Ryuk, NetWalker, Thanos, MountLocker, WastedLocker, Conti, Maze and Babuk strains.

"Criminals will always evolve their techniques to combine whatever tools enable them to best maximise their monetary gain with the minimum of complication and risk," noted Raj Samani, McAfee fellow and chief scientist. "We first saw them use ransomware to extract small payments from millions of individual victims. Today, we see RaaS supporting many players in these illicit schemes holding organisations hostage and extorting massive sums for the criminals." ■

Tunbridge Wells schools closed following data breach

Two Kent schools closed their gates and resorted to remote learning after hackers accessed their servers, stole data and encrypted pupil information.

Officials at the Skinners' Kent Academy and Skinners' Kent Primary School said they "cannot be sure" exactly what information hackers have access to, but

urged parents at the Tunbridge Wells schools to contact their banks to let them know that personal details could have been compromised.

Action Fraud and the National Cyber Security Centre are investigating.

The police and the trust's own data protection company are also carrying out inquires after the attack, which began June 2.

The Skinners' Kent Academy Trust said on its website that the hackers told them what information they have access to.

It said hackers not "appear" to have access to the School Information Management System, which is where personal records for pupils, students and staff are held.

"However, they have encrypted this data so that we no longer have access to it," the trust added.

A trust spokeswoman described the hackers as "sophisticated".

As staff no longer hold vital information on the pupils - including emergency contact details - the decision was taken to close the schools the following Monday.

The trust is now in the process of collecting all this data from parents again and the schools must also have their computers reconfigured so staff can access the resources required to teach.

A statement on the trust's website advised parents: "It would be very wise to let your bank know that your bank details may have been taken." ■

Hackers stole data from two schools in Tunbridge Wells

Lumen launches remote-work solution for organisations using Microsoft Teams

Lumen Technologies has introduced a remote-work solution for enterprises using Microsoft Teams.

The product, which is called Lumen Solutions for Microsoft Teams, is a managed, unified communications solution that leverages the Lumen platform to improve worker productivity, business agility and customer support.

"Few business Few businesses were ready for the challenges imposed by the global public health crisis," said Craig Richter, senior director of product management at Lumen. "With Lumen Solutions for Microsoft Teams, businesses can easily adopt 'work from anywhere' policies because Lumen takes on the burden of managing the complex calling and collaboration tools that make remote work possible. Delivering one of the world's most popular unified communications solutions over the Lumen platform gives customers access to a secure global network and a fast, secure platform for applications and data."

Since the start of the Covid-19 pandemic, unified communications and collaboration has become critical to maintaining business continuity and delivering superior customer experience because it integrates previously disparate voice, video and web services to provide a consistent and reliable user experience across a variety of devices.

Lumen said its new offering will see businesses benefit by combining the Microsoft Teams calling and collaboration tools "with Lumen expertise" in delivering a managed user experience that includes advanced reporting and analytics. ■

Proximity begins network enhancement

Proximity Data Centres has tasked Zayo Group with a "major network enhancement" for its edge colocation data centre at Chester Gates.

Zayo will deliver a high-capacity low latency fibre network to businesses located in the north west of England and a point of presence at the data centre, Proximity Edge 4.

The fibre network will allow for a range of high-speed services to be delivered, including ultrafast broadband connections. Furthermore, the PoP will offer low latency circuits to data centre hubs in Dublin and the USA, therefore allowing increased resilience for customer organisations requiring multiple connections to cloud services.

"Proximity's colocation data centre at Chester Gates is ideally located as a strategic point of presence for Zayo's high-capacity fibre network extending between Dublin and the north west of England," said Yannick Leboyer, chief operating officer of Zayo, Europe. "This collaboration supports the growing network demands and low latency needs of many more businesses and service providers across the north west region."

Proximity Data Centres

Proximity's growing network of UK regional edge data centres will expand to 18 sites "within the next 12 to 18 months", meaning the firm can offer UK-wide coverage with all its edge data centres selected for their proximity to major conurbation areas.

Built to tier 3 industry standards, all data centre grid electricity is sourced from 100% renewable providers. In addition, each data centre develops renewable energy solutions, including battery storage, solar and wind power wherever possible.

"Since acquiring two data centres in the region last year, which added a further six megawatts to our portfolio, we have had significant demand from businesses, applications developers and content delivery networks looking to use these high-quality facilities to move data and content closer to users," said John Hall, managing director of colocation for Proximity Data Centres.

"Collaborating with Zayo will allow us to offer enhanced services to these customers." ■

Iron Mountain sells five UK data centres for £128m

Iron Mountain has sold five of its London data centres to London-listed global asset manager Intermediate Capital Group (ICG) to the tune of £128m in a sale-leaseback deal.

The five facilities, all located in the greater London area, make up a total of 550,000 sq ft. of space and Iron Mountain will continue to be based in the different locations under an initial 12-year lease term. There is also an option to renew for up to an additional 20 years.

However, it remains unclear which facilities are included in the deal or if it includes data capacity it currently runs in Slough.

"This transaction is part of Iron Mountain's ongoing capital recycling programme and we expect to utilise the proceeds to reinvest in higher growth areas of the business," the company said in a statement.

Chad Brown, director at ICG added: "The Iron Mountain portfolio is a prime example of the mission critical real estate that ICG's sale and leaseback fund is seeking to invest in. This represents the fund's third transaction in 2021 and second transaction in the UK, following the 2.94 million sq ft forward funding of Jaguar Land Rover's new facility at Mercia Park earlier this year."

ICG said it has £500m available to invest in the UK and continental Europe. ■

Slough is the UK data centre hotspot

EDITORIAL:

Editor: Robert Shepherd
roberts@kadiumpublishing.com

Designer: Ian Curtis
ian@firstsightgraphics.com

Contributors: Steve Law, Kelvin Murray, Tony Gray, Peter Hudson, Sean Fitzgerald, Jeremy Wastie, Paul Ward, Robert Nitsch, Adrian Rowley, Anthony Senter, Shashi Kiran, Derek Granath, Karen Falcone, John Morrison, Brian Reed, Gemma Moore and Martin Hodgson

ADVERTISING & PRODUCTION:

Sales: Kathy Moynihan
kathym@kadiumpublishing.com

Production: Suzanne Thomas
suzannet@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

NHS Test and Trace strengthens cyber defences

NHS Test and Trace has selected Risk Ledger to manage cyber security risks in tits supply chain as a proactive measure to mitigate the increasing risks NHS and other critical national infrastructure organisations face from supply chain cyber-attacks.

The platform will give the service tools needed to manage cyber security risks in its supply chain at speed for a low per-supplier cost – apparently at least 60% cheaper than traditional solutions.

Cyber security risks in the supply chain can include third parties failing to encrypt sensitive data when it is being transferred. NHS Test and Trace, established to track and help prevent the spread of the Covid-19 virus in England, will take advantage of Risk Ledger's secure 'social network' allowing organisations to connect and share risk data securely, quickly, and easily. This gives organisations like NHS Test and Trace visibility of their supply chain and a comprehensive set of data to identify, measure and mitigate supply chain security risks at scale.

The move comes on the back of major supply chain cyber security breaches at Solarwinds and Microsoft in recent months, which have put the challenge of securing supply chains at the top of the agenda for organisations around the world.

"The government is working tirelessly to secure the nation online and grow the UK's £8.9bn cyber security industry as we build back better from the pandemic," said minister for digital infrastructure, Matt Warman. "We're helping SMEs develop innovative products and services and it's great to see Risk Ledger, one of the firms we've supported, win this contract to protect the Test and Trace system and support the national effort against coronavirus."

Risk Ledger's client base includes organisations like BAE AI, City of London Police, Telenor, Schroder's Personal Wealth and ASOS. The company's chief executive officer and co-founder Haydn Brooks described NHS Test and Trace as "essentially the biggest new start-up in the UK healthcare market". He added: "Healthcare organisations and their supply chains handle lots of highly sensitive data and have a high rate of data breaches. We have already seen during the COVID-19 pandemic that bad actors are actively targeting supply chains to access data and cause disruption." ■

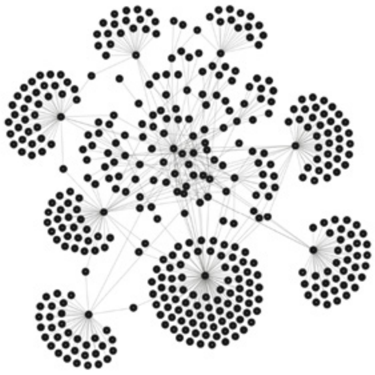

Subset of Risk Ledger supply chain network map

Scottish gov and agencies 'breached data protection rules almost 2,000 times'

The Scottish government and its agencies have breached data protection rules almost 2,000 times since the introduction of GDPR in 2018, a new report has claimed.

Disclosure Scotland, responsible for providing criminal record checks to employers and voluntary organisations, reported more than 200 data breaches.

The Scottish Prison Service was the worst offender and is responsible for more than 1,100 data protection breaches over the period, with more than 350-a-year since the introduction of the new laws three-years-ago.

Second worst was the Scottish government itself, which failed to report a serious data breach to the Information Commissioner's Office (ICO) within the 72 hours required by law.

Scottish Conservative chief whip Stephen Kerr said the number of breaches revealed "worrying weaknesses" in the Scottish government's security systems.

"Given that Scottish Government agencies store vast amounts of sensitive data, many people will be alarmed by these figures," he said. "Ministers must not take their eye off the ball when it comes to security-related issues. Urgent reassurances must be given that robust measures are in place to ensure the number of breaches is significantly reduced going forward and that all breaches are reported as quickly as possible."

Not all data breaches have to be reported to the ICO, with the seriousness of the breach determining whether a report should be made. ■

Government 'breached data protection rules'

FROM THE COVER

Getting to grips with NIST means better cyber security

The National Institute of Standards & Technology's (NIST) cybersecurity framework is often seen as a global standard for keeping businesses safe from cyber threats, but with the huge amount of useful information, it's easy to get lost in the detail. This summary will help get you started.

Identify and detect

The 'identify' and 'detect' elements of the NIST framework advise organisations to develop and implement effective ways to detect a cyber breach. This can take many forms, but scanning for breaches, anomalous behaviour and constantly checking data are important. If conducted manually by internal staff, this is time consuming, but automation can go a long way to lighten the load.

Powered by the latest AI and machine learning, a Security Information & Event Management (SIEM) platform can automate many processes and free up staff to investigate more serious events that require manual intervention. If you fall victim to a cyberattack, knowing about it quickly is essential in minimising the damage.

Protect

The 'protect' element of the NIST framework primarily cover prevention in the core areas of network, cloud and endpoint. The network perimeter is becoming ever more virtual, but that doesn't mean it's not important to protect with firewall, SD-WAN and DDoS protection solutions. Whether using public, private or hybrid cloud, clear lines of responsibility for data security and policy enforcement are critical. And finally, the endpoint or user is the most common breach vector so keeping users safe browsing the web, opening emails and downloading files is a key task.

Prevention also covers processes and people; according to a report created by the UK government, 48% of businesses have a basic cybersecurity skills gap. Consider outside help in the form of management service options, or a virtual Security Manager to act as an extension to your IT team.

The NIST framework advocates comprehensive awareness and training for all staff. Having systems in place to prevent a hacker accessing your network is no good if staff fall foul of a phishing e-mail with the same result.

Respond and recover

The 'respond' and 'recover' elements include response planning, mitigation and recovery activities to ensure continuous improvement. Start with an incident response plan covering key dependencies, backup and recovery solutions and any legal or regulatory requirements, to minimise damage and ensure you don't leave systems open to further attack.

A cyber breach can cause prolonged downtime - being able to restore systems quickly is essential to keep your business running and your customers happy. NIST can be an intimidating framework but focus on these core areas to get you started and move on from there and you will significantly strengthen your security posture.

By Steve Burden, Head of Security, Daisy Corporate Services, dcs.tech

Businesses in 551 towns and cities to get Openreach full fibre broadband

Openreach has revealed plans to provide full fibre broadband to 551 additional towns and cities across the country, covering some five million businesses and homes. The move is part of the company's £15bn programme to reach 25 million premises as it makes new gigabit cable technology available to 43,000 premises each week. The latest areas include Bournemouth, Kettering and Sunderland. "Our engineers and build partners are working flat-out to deliver this life-changing technology to rural, urban and suburban communities all over the country," Openreach chief executive Clive Selley.

Extreme Networks expands cloud footprint with new data centre

Extreme Networks has added a new regional data centre (RDC) in London, enabling customers to run its native cloud management platform, ExtremeCloud IQ on Microsoft Azure. The establishment of the London RDC means that Extreme's cloud footprint now extends to 17 data centres across the world. "This new location is helping us to further support cloud adoption for our ExtremeCloud IQ subscribers in region, delivering industry-leading information security and data protection with in-country data residency to stay compliant and advance their business," said John Morrison, SVP of international markets at Extreme Networks.

GTR to build UK's largest data centre

SEGRO has penned a deal with European build-to-suit data centre firm Global Technical Realty (GTR) to build its first UK-based facility and what will be the largest data centre campus in the UK hotspot of Slough. GTR, which is backed by global investment firm KKR, will take a 400,711 sq ft space for a 25-year term to operate bespoke data centres on behalf of global tech companies. SEGRO will make a bespoke site for GTR, which will be made up of three independent data centres capable of operating individually or as one interlinked campus. The project is expected to create 200 jobs during construction and a further 80 permanent roles on completion.

SonicWall introduces trio of new enterprise firewalls

SonicWall has introduced three new "high-performance" firewall models for enterprises and large organisations — NSa 4700, NSa 6700 and NSsp 13700 — designed to accelerate network throughput, stop advanced cyberattacks like ransomware and securely connect millions of users. The new appliances help enterprises keep pace with the speeds of their growing networks — all while drastically reducing total cost of ownership (TCO).

"The growing volume of ransomware attacks has enterprises and government agencies moving quickly to evaluate their mitigation capabilities and strengthen their security postures," said SonicWall president and chief executive officer Bill Conner. "The recent string of highly publicised cyberattacks has catapulted security to the top of the priority list.

Security pioneer McAfee found dead in prison cell

John McAfee, the creator of McAfee antivirus software, was found dead in a Barcelona prison on Wednesday June 23, hours after a Spanish court issued a preliminary ruling to extradite him to the US to face tax evasion charges. McAfee, who was 75, became famous in 1980s after releasing McAfee VirusScan that currently has nearly 500 million users worldwide. He was arrested in the Barcelona airport in October after being charged in Tennessee last year for tax evasion. McAfee was also indicted in New York in a cryptocurrency fraud case.

Schneider launches 'world first' liquid-cooled pre-fab facility

Schneider Electric has brought to market a world first liquid-cooled, pre-fabricated modular data centre. The EcoStruxure Modular Data Center, All-In-One Module is integrated by Avnet and contains chassis-level precision immersion cooling from Iceotope. The new prefabricated module will allow the most CPU and GPU-intensive high performance computing (HPC) edge applications to be deployed with greater reliability in harsh and remote environments. From industrial manufacturing and automotive sites, to telco, military, mining, oil and gas, the DC "enables real-time data to be processed faster with greater innovation, efficiency and lower latency", Schneider said.

EU rules UK data protection is 'adequate'

British data protection standards are "adequate", the EU ruled this week. However, Brussels warned that the decision could be revoked "immediately" if it sees a drop in UK standards. The decision is a boost for enterprises as failure to get a positive decision would have risked plunging British businesses into disarray, leaving numerous industries scrambling to set up more costly, bureaucratic alternatives to share data. The European Commission vice-president Věra Jourová said: "The UK has left the EU but today its legal regime of protecting personal data is as it was. Because of this, we are adopting these adequacy decisions today."

'Enterprise-run 5G networks open to risks', - GSMA report

Sizeable gaps in skills and tools mean enterprise-run 5G networks are exposed to security risks, according to new research. The latest *Securing 5G Era Private networks* report published by GSMA, found that 48% of surveyed operators see not having enough knowledge or tools to discover and solve security vulnerabilities as a key challenge. The survey further found that 51% of decision-makers worldwide said they prioritise IT and cloud vendor partnerships to improve network security.

Word on the web...

A successful cloud strategy is a secure one

By Rick Vanover, senior director of product strategy, Veeam

To read this and other opinions from industry luminaries, visit www.networkingplus.co.uk

Just like today's industrial leaders, Rajant's network is

Smart. Autonomous. Always moving.

Rajant Kinetic Mesh® is the only wireless network to power the non-stop performance of next-gen applications—from real-time monitoring to robotics and AI.

- Works peer-to-peer to maintain **hundreds of connections simultaneously for 'never break' mobility**
- Intelligently self-optimizes to **change in real-time, ensuring mission-critical reliability**
- The *only* network to enable **machine-to-machine communications required for autonomy**
- Provides **Industrial Wi-Fi for extended Wi-Fi connections** in challenging environments

IF IT'S **MOVING,** IT'S RAJANT.

*Industrial Wireless Networks **Unleashed.***

Why Carrier-Based LTE and Private LTE Just Aren't Enough for IIoT
RAJANT

Download our **"Why Carrier-based LTE and Private LTE Just Aren't Enough for IIoT"** white paper at **rajant.com/networkingplus**

RAJANT

Register for Networking+

Cybersecurity is not a one-stop-shop

Security should be more reinforced now than ever before, write Steve Law, CTO, Giacom and Kelvin Murray, threat researcher, Webroot

Despite lockdown restrictions easing, cybersecurity risks remain and are likely to grow as Covid-19 changes the working landscape. As indoor spaces begin to open in the next few months, employees will want to venture out to new spaces to work, such as coffee shops and internet cafes – but working on open networks and personal devices creates unlocked gateways for cyberattacks to take place. Since this hybrid and remote way of working looks like it's here to stay, businesses must ensure they have the right infrastructure in place to combat any cyber threats

For instance, research by the National Cyber Security Centre shows that there has been a rise in Covid-19 related cyberattacks over the past year, with more than one in four UK hacks being related to the pandemic. This trend is not likely to ease up any time soon either. And, going forward, hackers could take advantage of excited travellers waiting to book their next holiday once the travel ban is lifted, deploying fake travel websites, for example.

Aside from the bad actors in this wider scenario, part of the problem here is that many IT teams are not making use of a holistic and layered approach to security and data recovery; which can lead to damaging consequences as data is stolen from organisations. Such issues continue to resonate strongly across businesses of all sizes, who will, therefore, turn to their MSPs for a solution.

The importance of a layered approach

Cybersecurity is not a one-stop-shop. A number of solutions are required to ensure maximum effect. They includes a layered combination of DNS networking, secure endpoint connections and an educated and empowered human workforce.

The need for DNS security cannot be ignored, especially with the rise of remote workforces, in order to monitor and manage internet access policies, as well as reduce malware. DNS is frequently targeted by bad actors, and so DNS-layer protection is now increasingly regarded as an essential security control – providing an added layer of protection between a user and the internet by blocking malicious websites and filtering out unwanted material.

Similarly, endpoint protection solutions prevent file-based malware, detect and block malicious internal and external activity, and respond to security alerts in real-time. Webroot® Business Endpoint Protection, for example, harnesses the power of cloud computing and real-time machine learning to monitor and adapt individual endpoint defences to the unique threats that users face.

However, these innovative tools and solutions cannot be implemented without educating users and embedding a cyber security-aware culture throughout the workforce. Humans are often the weakest link in cybersecurity, with 90% of data breaches occurring due to human error. So, by offering the right training and resources, businesses can help their employees increase their cyber resilience and position themselves strongly on the front line of defence. This combination is crucial to ensure the right digital solutions are in place – as well as increasing workforces' understanding of the critical role they play in keeping the organisation safe. In turn, these security needs provide various monetisation opportunities for the channel as more businesses require the right blend of technology and education to enable employees to be secure.

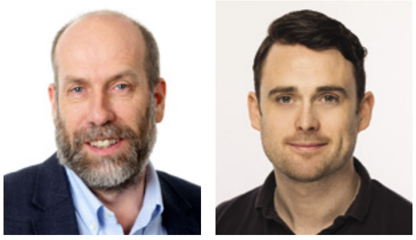

Steve Law, CTO, Giacom and Kelvin Murray, Threat Researcher, Webroot

The channel's role

Businesses, particularly SMBs, will look to MSPs to protect their businesses and help them achieve cyber resilience. This creates a unique and valuable opportunity for MSPs to guide customers through their cybersecurity journeys, providing them with the right tools and data protection solutions to get the most out of their employees' home working environments in the most secure ways. Just as importantly, MSPs need to take responsibility for educating their own teams and clients. This includes delivering additional training modules around online safety through ongoing security awareness training, as well as endpoint protection and anything else that is required to enhance cyber resilience.

Moreover, cyber resilience solutions and packages can be custom-built and personalised to fit the needs of the customer, including endpoint protection, ongoing end-user training, threat intelligence, and backup and recovery. With the right tools in place to grow and automate various services – complemented by technical, organisational and personal support – channel partners will then have the keys to success to develop new revenue streams too.

Conclusion

Hackers are more innovative than ever before, and in order to combat increasing threats, businesses need to stay one step ahead. Companies must continue to account for the new realities of remote work and distracted workforces, and they must reinforce to employees that cyber resilience isn't just the job of IT teams – it's a responsibility that everyone shares. By taking a multi-layered approach to cybersecurity, businesses can develop a holistic view of their defence strategy, accounting for the multitude of vectors by which modern malware and threats are delivered. Within this evolving cybersecurity landscape, it's essential for SMBs to find an MSP partner that offers a varied portfolio of security offerings and training, as well as the knowledge and support, to keep their business data, workforces and network secure. ■

Separating the wheat from the chaff

Martin Hodgson, head of UK & Ireland, Paessler AG

Remote working – considered by some a "luxury" but now very much normality – has taught us many things over a large part of this year. With face-to-face interaction limited, the amount of messages and notifications we receive on a daily – even hourly – basis has increased. As a result, many of us feel a sense of bombardment, without the ability to navigate what priorities are present and what activity can take a backseat. Sound familiar?

Let's now scale this up to think about what those managing our corporate networks whilst at home must be dealing with. During 'normal' circumstances enterprise IT environments are complex covering hundreds and thousands of devices and applications from many different vendors. In this environment, alert noise sits as one of the biggest issues that IT teams face.

This happens when you're monitoring your infrastructure, network, storage, cloud services and other elements of your IT, and generating alerts and notifications for failures or impending issues. Too much alert noise makes it downright difficult to identify serious problems, and it might even mean ignoring alerts and missing what really matters. Your monitoring efforts are compromised, and the quality of your service goes down.

How to reduce alert noise

IT teams are key to ensuring businesses continuity in the current climate, so reducing alert noise needs to be a priority. This takes a combination of careful, strategic planning paired with the right monitoring tool. The below four methods should give guidance on how to separate the wheat from the chaff.

1) Consolidate your monitoring into one tool

Streamlining the way you are alerted by having all devices feeding into one tool is the first step to quietening the noise. This ensures that when alerts come through, you only have to refer to one tool to find the underlying problem. Additionally, because tools handle alerting and notifications differently, a single tool means that you can apply the same philosophy across the board.

2) Effectively set your thresholds

Alerts are based on thresholds. For example you'll get an alert when a device overheats past the set threshold or when storage is lower than the set requirement. To avoid having numerous incorrect alerts being triggered due to standards being incorrectly set, you'll need to review the thresholds as part of a good alert management process. Set them too low, and you'll get inundated with alerts; set them too high, and you won't get notified when there's an issue until it's already too late.

Alongside this, when managing multiple devices it is crucial to have a monitoring solution that offers automation and other mechanisms like inheriting thresholds for groups of devices.

3) Distribute alerts to the relevant teams

Thirdly, you need to have a monitoring tool with comprehensive rights and roles functionality. This way you can easily create roles and responsibilities for specific teams (or even individuals), then filter alerts accordingly.

For your monitoring concept, define the user groups according to the areas that they focus on. Then, you define notifications for failures in those areas to go to the specific teams that need to know. For example, you might have an IT team that handles your online store, and another team that handles the email services. In this example, you would configure that the team handling the online store only receives alerts relevant to that area, and the same for the team handling the email services. This way, alerts get sent only to the relevant teams.

4) Only send high-level alerts to senior management

A key thing to remember is that not everyone in your organisation needs to know what's going on behind the scenes of your infrastructure. Often decision makers, management, and other business stakeholders only need to know the health of the network at a very high level. Organising your infrastructure into IT services according to business processes can do this. For example, if there's a service-critical problem and you have the right alerts in place, an alert can be sent to relevant management members or stakeholders.

Noise reduction for optimised business functions

So, in conclusion, reducing alert noise takes time and consideration, yet it will allow both the IT teams and the relevant business leaders know when immediate action needs to be carried out and when simple solutions can be implemented. By having just one reliable monitoring tool in place, enterprises of any size can streamline the alert types, set standards and delegate responsibility. What's more, it means they only need to focus on what is critical, to ensure the smooth functionality of the network. ■

Is there a 'best' critical comms solution or is it always a combination?

Networking+ **caught up with the biggest players in the critical communications space to get their views on the current critical comms space and the technology that makes it possible**

What are key critical comms issues, considerations for the UK marketplace?

Tony Gray, chief executive, TCCA: The UK is currently working to replace the Airwave TETRA network used by all the emergency series. Whilst still functional and still very good for voice communications, the decision to move away from a narrowband network to a standards-based 4G LTE network is aimed at allowing the seamless integration of voice and broadband data, greater choice of devices and a clear roadmap forward as technology evolves. A system more aligned with the MNO model should allow for greater competition in the provision of the network aspects, by separating it away from the functionality provided to run over the top. By splitting the functionality and allowing competition the costs can be reduced.

The issues from a future perspective are what proven standards-based solutions exist and currently this has yet to be fully demonstrated in a live operational environment. The parts for a full solution exist, but the ability to bring them together to form a single system that provides all the functionality that current TETRA systems offer has yet to be seen. The other main issue is the need to try and bring the various parts of a new system together at the same point in time. This is made far harder as the parts, such as devices, follow much more closely the refresh cycles of modern smartphones, so have a lifespan much more in the three-four year period, rather than the seven-nine year lifespan of a TETRA device. This means that the device will need to change during development, testing and the transition to live service, with every change of device requiring full regression testing, using up significant parts of the device lifespan. This same constant refresh cycle will be seen across many of the parts of the overall system. On the upside this constant refresh cycle brings greater functionality with every release.

Sean Fitzgerald, head of EMEA solutions marketing at Motorola Solutions: Even though it's a mature market, critical communications technology is continually evolving. The UK mission-critical market is constantly adapting to changing demands and challenges, such as operations becoming increasingly data-driven. Over the past year, the pandemic also accelerated the need for solutions that facilitate remote working as social distancing regulations prompted closure of many workspaces. For many, this meant a greater focus on collaboration and productivity, including applications that make individual's and team's workflows more efficient, and highlighted the need for simple communication using voice, video and data.

This, in turn, led to wider adoption of cloud-based solutions for data sharing and storage. A collaborative ecosystem enables more seamless end-to-end workflows, helping teams and individuals across multiple sites communicate with voice and access data. Reliable, accessible solutions are essential features for sharing secure, robust mission-critical communications.

Cellular or TETRA – which best fits specific user groups?

Peter Hudson, chief technology officer, Sepura: TETRA networks are specialist networks delivering mission critical communications services, with specific feature sets that have evolved and improved over time. These requirements are based on extensive user interactions, understanding how people interface with technology whilst undertaking their role.

TETRA networks are designed to provide a high level of inherent resilience and redundancy in their architectures. The networks are dimensioned to provide a specific grade of service at peak load to the user groups they serve; if peak demand is overreached, calls are queued, rather than dropped.

Cellular networks are designed and operated to meet a different set of criteria, user needs and operational efficiency. These demands are often conflicting to those of a mission critical user. Cellular networks can be configured to fit the needs of critical users, but this can be a tradeoff to commercial use. Private or national networks can be configured to the needs of mission critical demands with full end to end integration, appropriate terminals or devices and compliance to appropriate (3GPP) interface protocols and standards.

In practice the split between cellular and TETRA is not so black and white. Hybrid or mixed networks provide the best of both worlds, delivering what they are best designed for. Broadband or LTE networks are inherently good at delivering data services, where the response and reliability are not mission critical, whereas TETRA is proven to provide reliable mission critical voice communications. Use and integration of both networks delivers the gold standard service with the opportunity to explore and benefit from a whole range of new data applications. Broadband technology is complementary to narrowband - it does not necessarily replace it (certainly in the near to mid-term) but enhances what can be offered to users.

The choice between cellular and TETRA is based upon user needs.

Tony Gray, chief executive, TCCA: As the question suggests, it is very much a matter of the particular types of users and their needs. Cellular is designed for consumer subscribers, supporting one-to-one calls with relatively slow set up time, can have coverage issues in some cases and may be unreliable at times of peak demand or overload. Communications for TETRA users are mission critical and may even be lifesaving. If the user need is for reliable, resilient, secure group communications with sub-0.5 second call set up and all the other features well known, proven and trusted by users, then the choice has to be TETRA.

Which three requirements must mission critical comms meet?

Peter Hudson, chief technology officer, Sepura: Reliable coverage – the network must provide reliable communications, always available in all circumstances. Quality of service and performance KPIs must be sustained to ensure successful operations, both within the network coverage and when network is not available. Direct mode operation is a key operating requirement for most mission critical users who need to be able to rely on communications however the network status changes, planned or unexpected.

Robust design – the devices must be sufficiently robust to keep working in adverse weather conditions – whether hot, cold, or wet – or adverse environmental conditions, for example in dusty, salty or dirty locations. The devices must also be tough enough to stand repeated cycles of deployment with varied users, withstand heavy kicks, drops and intensive use.

Practicable usability -– audio must be intelligible, loud and clear so that information and instructions are heard first time, every time. This includes whilst operating in noisy, dangerous or busy environments. This intelligibility is enhanced by a portfolio of mission critical quality proven, accessories allowing users to use devices in a manner that supports their daily operational needs, for example eyes free usage and easy to access using gloves and other protective equipment.

Jeremy Wastie, head of public sector sales, MLL Telecom

1. Capacity – reliable symmetric performance – more and more data is going "upstream" rather than simply being down loaded from the cloud.
2. Resilience - high SLAs with resilient fall back – a 10 minute outage can be business critical or even life threatening in Critical National Infrastructure. Networks need to be designed with multiple layers.
3. Secure – end to end security of data.

Tony Gray, chief executive, TCCA: There are more than three, but top of the list typically come reliability, resilience and security, wherever and whenever users need to operate. Mission critical communication users often operate in dangerous situations in challenging environments, so their communication services need to be available always and everywhere – it is impossible to predict when and where incidents take place. Good geographical radio coverage (both on network and off network) is essential for high service availability. The quality of communication services needs to match user requirements, seamlessly supporting users' operational processes and supporting user needs depending on the incident and the operational processes.

There are some aspects of spectrum management where government can support the deployment of a 4G LTE network, for example ensuring some spectrum is authorised for use above 500ft, to assist with the deployment of Air to Ground networks.

What's the biggest problem facing critical communications today?

Jeremy Wastie, head of public sector sales, MLL Telecom: Consistent high capacity services across the UK. This is being addressed by the emphasise on full fibre networks across the UK both from the independent new players but also from the established telcos who are investing their own and government funds in developing full fibre service and switching off the long-serving but historic copper networks. Closely behind we will see the growth in 5G services from the mobile operators, but these will rely on those full fibre networks to deliver the necessary backhauls for 5G services to really fly.

Tony Gray, chief executive, TCCA: It's not a problem unique to critical communications, but budgets and priorities are a universal issue faced by many users in all walks of life who need investment in systems and technology to make their lives and work more efficient, effective and safe. Although the most important area of communications, the mission critical world is very small compared to the consumer market. Business cases based on volume cannot compete with the consumer landscape, so while there is a huge amount of innovation in our sector, only a fraction makes it to market as truly mission critical products and services. Resource is also an issue – many of the key developments in critical communications have been and are driven by committed volunteers who work for the good of society and not for the profit.

Is unified critical comms something you recommend or does it but too much additional pressure on a network?

Jeremy Wastie, head of public sector sales, MLL Telecom:

1. Standard voice services e.g. MS Teams are now becoming a de facto cloud based service and can always be backed up by mobile services. Provided it is built into network design.
2. Contact Centre and Emergency Command and control services should remain "standalone".

Peter Hudson, chief technology officer, Sepura: No one technology will provide the answer to all needs, either technically or economically. To provide the range and capability of services that we can see being demanded, a range of technologies will be required to achieve this; essentially creating a network of networks.

To reduce complexity in cost and operation it is essential to define carefully which elements need to communicate with each other, rather than a default 'everything must interoperate with everything' approach. A unified critical comms approach also provides a pathway to adopting new technologies or developments as they become available, ensuring the cost of change is minimised.

It is essential that stakeholders have a sure understanding of security requirements when adopting a hybrid solution. Only in this way will they be able to ensure that any end to end communication, in particular cross technology communications, is fully secure and safe from attack.

Sean Fitzgerald, head of EMEA solutions marketing at Motorola Solutions: The reliability of radio remains a key for critical communications. There is, however, significant growth in the importance of other data, including video, and subsequently a greater need for unified critical communications. Many businesses also require collaboration with other devices that their workers carry, including LTE devices and body-worn video cameras, without compromising the security or reliability of their radio voice communications.

As the use of video in critical communications becomes more important across a range of sectors, devices and software have evolved to make it simple for users of different devices to share data over networks, and store video and images in cloud-based solutions for easy, remote access.

A trusted partner for unified critical communications understands the challenges their user faces, developing solutions collaboratively with the customer that can meet and support their needs. An approach enabled by a common ecosystem of hardware and software, such as the WAVE PTX application, allows users to share voice, video, images and other data with MOTOTRBO or TETRA two-way radios, plus broadband and LTE devices at the push of a button, working seamlessly across networks for reliable communications.

What is the most unusual critical comms application you're experienced?

Paul Ward, international sales director, Etelm: Etelm was involved in a project for biometric monitoring of military personnel. This was following a number of incidents where armed forces personnel died or became seriously unwell whilst on exercises or during remote tactical deployments. The system involved monitoring all key cardio and biometric information of troops in the field, using Bluetooth data to the personal mobile device which was then transmitted in real-time over our critical communications network. Even core body temperature could be monitored using a capsule that could be safely swallowed, sending Bluetooth data instantly over the network to the command centre. Alarms would be triggered based on specified parameters so any health concerns could be responded to quickly.

Robert Nitsch, Frequentis vice president public safety: Unusual applications are in the eye of the beholder. By definition, emergencies are unexpected events that often require immediate action. However, there are also planned events that put an unusual level of strain on resources or accentuate critical communication needs; for example events such as the Olympics or the G7 summits, both of which Frequentis has been able to support with its critical communication solutions, already implemented around the globe.

The need to temporarily scale up operations, to handle the surge in communication demand, to manage such event, and to be able to seamlessly patch different communication technologies (analogue and digital radio, telephony), to safely reach and coordinate emergency personnel to save lives, protect property and maintain public order is vital. With a web-based front-end our 3020 LifeX solution is capable of providing access to an operator working position from anywhere, meaning dispatchers can manage emergency incidents from any mobile device or location. This was also very beneficial during the pandemic where extra control room resources were needed to manage increased caller demand.

In what situation do you think the enterprise user should consider a critical comms solution?

Paul Ward, international sales director, Etelm: There is a big demand from private critical communications users for broadband mobile data services. This can never be achieved using traditional narrowband PMR/LMR services – TETRA can offer secure voice and low speed data but for users wishing to use high speed mobile data applications LTE is definitely the best choice. As LTE voice standards evolve the option to utilise the best features of TETRA for mission critical voice and LTE for high speed data services is ideal as the industry moves towards full mission critical LTE.

Robert Nitsch, Frequentis vice president public safety: There are many enterprise customers with critical communication needs, for example airports, and utilities There is no reason why they should not consider the same solutions as Emergency Services, especially when they are available as a software as a service (SaaS) model. Often it is just the assumption that Emergency Services systems are bespoke and therefore too expensive.

The Frequentis multimedia communication platform 3020 LifeX, for telephony and radio dispatching, is unique and was especially designed and built for current and future communication demands. It is able to support all types of communication methods and display those in an easy to use irrespective of new media and associated communication features, such as ESN (MCX) and NG999, nation-wide communication solutions, integration of AI services or SaaS. ■

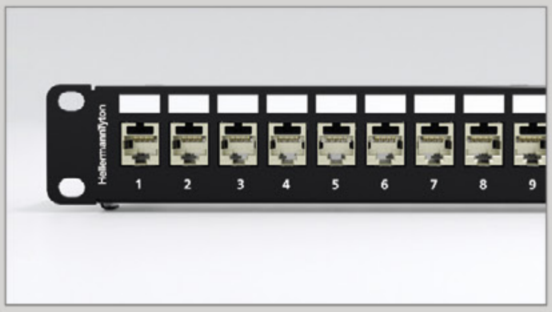
HellermannTyton
1
2
3
4
5
6
7
8
9

HellermannTyton
1
2
3
4
5
6
9
10
11
12
13
14
15
6A
HT

SD-WAN products focus

The SD-WAN market continues to grow unabated. So, for those looking to take the plunge, or find a new provider, here's what's out there

There's been a lot of buzz around SD-WAN for some time now and the market is exploding alongside digital transformation.

What's more, there are different reasons for this. One is the fact it forms key transition for enterprises, particularly as they expand their capacity for enabling work across additional distributed location.

The UK has undergone a massive transition in the way people go about their everyday lives. Organisations that have embraced the nine-to-five working culture of the past several decades have found themselves rushing to adopt new ways of working in order to keep their heads above water in what has been a turbulent 16 or so months, courtesy of the Covid-19 pandemic. In other words, things have been catalysed by the influx of remote working that the vast majority of us had no choice but to embrace at the beginning of 2020.

To that end, SD-WAN is capable of forming the backbone of your organisation's digital transformation strategy. The increasing adoption of cloud-based applications will soon necessitate a more advanced, intelligent networking fabric than MPLS. SD-WAN's features provide the intelligent functionality that optimises an enterprise's network and resolves remote working challenges.

Time then to take a close look at what products are out there and the companies behind them.

First up is Aryaka. To date, the company has been known for its global managed SD-WAN offering, according to Shashi Kiran, the company's chief marketing officer. "This integrates connectivity across our private SLA-driven core, DIA, and MPLS, as well as regional multi-cloud access, security, optimisation, and visibility," adds Kiran. "Our end-to-end SLAs, application performance, and customer care are considered to be unequalled in the industry."

Kiran is effusive about the recent acquisition of Secucloud, a Germany-based SASE company "that will accelerate our security roadmap with a goal of offering choice to our customers, either via our organic SASE capabilities delivered in conjunction with our SD-WAN offer, or those of our security partners" that include Check Point Networks and Palo Alto Networks.

As Kiran opines, SD-WAN, extended to include SASE and offered as a managed service, meets the transformation requirements of enterprises by delivering flexibility, scalability, simplicity of operation and TCO advantages. "It enables them to better compete in a highly competitive and dynamic environment," he adds.

Aryaka's customer base, thus far, has been multinationals that have application performance and security requirements across multiple regions. "Our fully managed service including an SLA-driven core delivers the user experience they require to ensure productivity for their global employees. We are now building on this offer with services aimed at regional deployments," says Kiran.

The eponymous SDWAN Solutions is run by chief executive officer Anthony Senter. The company spent 18 months working with one of the world's largest hardware manufacturers to develop its own range of universal CPE. Its desktop Vena (Virtual edge network appliance) and rack mounted Vecta (virtual edge core termination appliance) devices service requirements from 20Mbps up to 5Gbps throughput and allow it to run most SD-WAN software and security software (via VNF capability) on a single device. "This also means that our solutions are not restrictive as we can replace the SD-WAN vendor and /or the security vendor without having to replace hardware, giving customers best of breed tailored solutions, increased performance and capability," says Senter. "Our solution OPEX pricing model guarantees no huge upfront costs for customers which makes post pandemic network and router replacement projects budget friendly. We can also deliver and install our hardware in 190 countries with four hour or NDB RMA."

Senter says the company has a few different pricing models when selling via the Channel and direct, and these also vary depending on quantity ordered as well as the installation method – i.e. fully managed, zero touch deployment or DIY. "What we can say is that our entry level Vena-I device is comparable to many provider's medium sized devices,but are priced to compete against those providers entry level device," he says.

Aruba has been in the SD-WAN game for some time now. It offers the latest version of the Aruba EdgeConnect SD-

WAN edge platform software, acquired with Silver Peak, is 8.3. Derek Granath, senior director of product marketing at the company, says Aruba EdgeConnect helps geographically distributed organisations of all sizes and in all verticals connect users to applications with the highest quality of experience wherever they reside. Its current customers range from legal firms to retailers, construction companies to financial services firms – with recent examples including Lush and Barrett Steel.

"At the highest level, enterprises with geographically distributed locations and that are using applications hosted in the cloud (SaaS, IaaS and PaaS) benefit most from the efficiencies, business agility (i.e. ability to accelerate deployment of new sites and applications) and application performance and availability improvements that the Aruba EdgeConnect SD-WAN edge platform delivers (i.e. the ability to always connect to applications with the highest quality of experience)," he says. "We also have a number of CSPs that offer the Aruba EdgeConnect platform as a managed SD-WAN service to their market segments.

In terms of user groups, "Aruba participates in the MEF and ONUG to continue to advance the innovation and adoption of SD-WAN technologies. SASE and security are big themes/topic for both, and it is really useful for us to get multiple perspectives from different user groups", says Granath.

Hot off the press, Vapour has launched a Fortinet-backed SD-WAN solution for hybrid workers.

Deployed as a software licence – with three levels of support available on a 'price per user per month' basis – the solution facilitates endpoint visibility and management, secure remote access for employees irrespective of location, and automated threat intelligence and response.

Mindful that IT leaders are seeking zero-touch provisioning, this technology is even simpler to roll out, as no hardware installation is required. The service has been 12 months in the planning.

"We all know that the traditional 9-5 is long gone, with employees seeking increasing flexibility from their working day –especially when it comes to location.," says Vapour's chief executive officer, Tim Mercer. "Many people are keen to remain at home, some are hungry to return to the workplace, and others want the fluidity to decide from one day to the next."

"Hybrid working is definitely the future, but this isn't the easiest to manage from a network security perspective. This is exactly why we've launched this solution – it's powerful technology designed to liberate businesses, as they move forward in a post-Covid world."

Juniper needs no introduction. As Karen Falcone, senior director of SD-WAN/session smart routing product marketing explains, traditional SD-WAN offers a number of key benefits enabling workforces to adapt to the remote and hybrid working model that was thrust upon us. However, she argues that Juniper's SD-WAN solution takes this to the next level by combining its "unique session based SD-WAN" with the power of automation and AI with its Mist AI.

"The Juniper AI-driven SD-WAN combines our Session Smart SD-WAN solution with insights from Mist AI to meet the agile needs of today's enterprises, and to ensure the user experience comes first," she says. "Whilst traditional SD-WAN solutions might report an 'up' link as operational, it can lack the additive intelligence to infer adequate network performance and overall user experience. Our session-based approach provides unprecedented granular visibility, insights and control, transcending the network layer to a user and application level."

Falcone adds that unlike traditional tunnel-based SD-WANs, Session Smart is inherently built on zero trust and validates all sessions with a global user and application policy. "When combined with Mist AI, we leverage session insights to understand the issues that are affecting network clients in real-time," she adds. "Customisable service-level expectations (SLEs), proactively detect and intercept anomalies and issues as they occur and offer intelligent and considered network recommendations through our Marvis virtual assistant in natural language." ■

Derek Granath, Aruba

Shashi Kiran, Aryaka

How do SLAs work?

"I think the industry has moved on from SLAs – at SDWAN Solutions we pass each connectivity providers SLAs through to the customer, but when it comes to a solution as a whole, we would much rather design the solution correctly to eliminate single points of failure and build in redundancy and resilience, rather than offering our customers a £30 rebate for an eight-hour outage. It scares me when I see providers offering 100% SLA's based on fundamentally flawed designs and installations – that's more of a marketing approach rather than an actual business benefit."

Anthony Senter, CEO, SDWAN Solutions

SLAs have created a great deal of confusion and misplaced expectations, and there are vast differences between 'marketing' SLAs and those that have business and technology 'teeth.' One critical aspect is whether the SLA covers just a portion of the network, or if it is truly end-to-end. Also, if a vendor or provider is claiming 5-9s, for example, whether they have the instrumentation and verification tools to keep to their commitments. This includes the different WAN connectivity options, the first, middle, and last miles, and the visibility into both the overlay and underlays if applicable. All of this combined is the 'transparency' that is many times missing in various SLA claims.

Shashi Kiran, CMO, Aryaka

Are there any security concerns and how do you deal with them?

Aruba EdgeConnect includes a comprehensive suite of unified security functions that secure the data plane and the management and control planes. Data plane security includes 256-bit AES encrypted paths within the SD-WAN fabric, branch security features including a stateful zone-based firewall, dynamic segmentation, DDoS detection and more. First packet application identification enables granular traffic steering to enforce business-driven security policies across all sites on the network.

Derek Granath, senior director of product marketing for Aruba.

In the past, IT has sometimes looked at SD-WAN and security as two separate silos, as opposed to a converged solution, resulting in increased complexity and risk. This has changed over the past few years with SD-WAN edge devices that integrate security or cloud-based offers that interwork with the SD-WAN edge. More recently, the advent of SASE has helped to bridge any gaps between the networking and security teams. However, there is still a major burden on the IT team to properly plan and manage this combined architecture, including both on-premises and cloud-based security where required. A way out of this is via a managed service where complexity and risk are removed from the IT team as part of a fully managed SD-WAN and SASE offer.

Shashi Kiran, CMO, Aryaka

Our biggest concern is the lack of current security in most customer networks. Cyber criminals are launching generation 5 and 6 attacks against companies but 95% of businesses are only protected against generation 2 or 3 attacks. Cybercrime increased 38% during the pandemic. Add in WFH and BYOD and companies are more exposed than they think. SDWAN Solutions partners with the top 3 security vendors to offer secure edge (SASE) and cloud security seamlessly integrated into our SD-WAN network solutions to protect our customers. You can do a quick online Check Point CheckMe proactive assessment that identifies security risks on your network, endpoint, cloud and mobile environments, on our website

Anthony Senter, CEO, SDWAN Solutions

What is the 'hybrid cloud visibility gap' and why does it matter?

Adrian Rowley, senior director EMEA at Gigamon

Adrian Rowley offers insight into the pervasive issue of the hybrid cloud visibility gap, and explains why observability is key for cloud migration, security and customer experience.

The 'hybrid-cloud visibility gap' is a direct consequence of a period of accelerated digital transformation and the rapid, sometimes unplanned and unprepared for, shift to the cloud. The IT landscape that we knew before the COVID-19 pandemic has altered significantly, meaning that many digital tools and traditional processes have become practically obsolete. While network monitoring remains a top priority, full visibility into all data-in-motion has often been lost as organisations have rushed to implement multi and hybrid cloud infrastructure alongside their on-premises systems, yet many legacy tools do not stretch seamlessly to the cloud environment. In fact, visibility has often been completely restricted, so much so that each cloud embodies its own island of visibility. This creates a 'gap' that needs to be bridged in order for NetOps teams to once again achieve a unified view of their network and ensure that potential cyberthreats are detected and data processing is optimised.

During a time in which businesses must innovate to survive, cloud infrastructure has become integral for success and will continue to be important as the hybrid workforce looks set to stay. Yet with 90% of organisations stating their cloud usage has become more than they initially planned for, it seems that not many businesses are prepared for securing and scaling their cloud operations, and some NetOps teams may feel out of their depth. The hybrid cloud specifically, a combination of public and private clouds, will continue to grow and is currently already being deployed by 82% of IT professionals. Yet, if visibility is not prioritised, it could cause a number of worrying issues for those hoping to drive recovery in the post-COVID world.

According to a recent poll, 40% of respondents claim that a lack of visibility is one of their main concerns when migrating to the cloud. It is impossible to manage what you cannot see, and if visibility is 'clouded', migration will become far more costly and complex. In fact, digital initiatives may fail completely if they become over-complicated and if there is not sufficient visibility to successfully rebuild workloads within the new environment. A visibility gap across the hybrid cloud network means that security, compliance and performance issues are likely to arise and digital transformation could be scuppered completely. What's more, as additional teams, tools and agents are feeding into the complex migration process, there will also be an unnecessary increase in network traffic, which has expensive repercussions for bandwidth and CPU capacity. Unified visibility is key to reducing the complexity of the hybrid environment, and a single pane of glass view into traffic can mean that the movement of data is optimised and therefore costs are reduced.

Security within the cloud is one of the biggest concerns within the IT environment – with 81% of organisations viewing cloud security as a challenge. A gap in cloud visibility will inevitably mean that a business is more vulnerable to cyberattacks, as it then becomes impossible to monitor all traffic, or detect all threats. Without full visibility, SecOps teams must turn to less reliable sources of information, like application logs or trace files, yet this level of insight will not protect a cloud environment from the increasing number of attacks. It is important to remember that the instrumentation of logs and applications can vary developer to developer. For example, it is common to see logging levels minimised by CloudOps teams to improve performance, but compromise on security. Therefore, to better ensure security and compliance within the hybrid cloud, all data-in-motion should be visible, including east-west traffic from containers and unmanaged devices. Only then can SecOps teams be confident with their cloud security posture.

If cloud security is compromised, digital transformation becomes overly complex, networks remain far from optimised and customer experience will inevitably decline. When a NetOps team is unable to glean a clear view into network traffic, application problems can go unnoticed and network bottlenecks can build up. While customers are already enduring a slow digital experience, IT teams will then struggle to remedy issues quickly if the hybrid cloud visibility gap remains un-bridged. One or two disgruntled customers is an issue many organisations must tackle, yet a large amount of disruption on your network could have more far-reaching consequences and cause significant damage to revenue, as the end-user looks to your competitors. However, if visibility once again becomes a priority, troubleshooting issues can be a more streamlined process and NetOps teams can filter out low-risk, duplicate or irrelevant data to free up bandwidth, speed up the network and ensure a better digital experience for customers.

It is clear that it is not simply IT professionals that feel the consequences of the hybrid cloud visibility gap, but that the whole business will suffer its repercussions if left unmanaged for too long. A unified view into all data across your network is essential to secure your workforce and satisfy your customers, whether that is simply between a public and private cloud, or spanning a number of cloud and on-premises environments. Without this visibility, maintaining a high-functioning and secure digital environment becomes more challenging than ever before. ■

Building the cities of the future

Cities around the globe are increasingly looking to digitalise their infrastructures in an attempt to improve the overall quality of life for their citizens. John Morrison, SVP of international markets, Extreme Networks, highlights the most important factors to consider when creating a smart city strategy and touches on ways cities can remain secure as they digitally transform

Covid-19 exposed the shortcomings of major cities and public service providers that did not have a clear digital strategy. As a result, a study by EY revealed that 62% of HHS (Health and Human Service) organisations rightly increased the use of digital during the pandemic. The past year has exposed an urgent need to digitally transform in order to improve public services and prevent such a catastrophic event from damaging urban areas in the same way again.

The pandemic also revealed that cities should and could be better at harnessing the power of technology and data analytics to improve operational efficiency. In the US, for example, Chicago has used anonymised mobile phone data to analyse travel patterns and track whether people were self-isolating or not. What's more, globally, manufacturers transformed their production lines to produce medical supplies and equipment, underscoring the importance of partnerships between infrastructure, technology and mobility experts.

From improving public safety to enhancing social connectedness, and even helping to create cleaner and more sustainable environments, the smart city vision promises to solve the challenge many cities are facing today. Making a city 'smarter' improves quality of life while reducing costs and generating economic growth.

Smart cities utilise technology to effectively manage assets and resources that in turn improve the quality of services provided to residents.

Technology has transformed over the past few decades. Today, everything in our homes, or on the street, can be wirelessly connected and thankfully, dial-up internet has become a distant memory. Consistent connectivity is important for any smart city to function and so a strong infrastructure is needed to make sure that a vast ecosystem of devices, from anyone and anywhere in a city, can remain truly connected.

A fundamental component of a smart city is Internet-of-Things (IoT) technology. An IoT-enabled city can divert traffic to avoid congestion in real-time, detect faults in key infrastructure such as street lights or road signs and can even monitor energy usage across a city to help manage and reduce levels of pollution. Smart cities that utilise IoT-enabled technologies will be able to not only maximise operational efficiency, but enhance residents' quality of life.

But inevitably as the number of connected mobile and IoT devices continues to grow, so does the number of potential vulnerabilities to be exploited, increasing the risk of cyberattacks. With the increased number of IoT devices, the potential attack surface significantly expands, creating entirely new vulnerabilities to protect against.

The total installed base of Internet of Things (IoT) connected devices worldwide is projected to amount to 30.9 billion units by 2025, a steep increase from the 13.8 billion units that are expected in 2021. But although companies are producing connected devices at pace, gaping holes are potentially left in the form of weak security. Many firms are developing IoT firmware with open source components in a rush to release their products into the market. But this can cause serious problems.

As these devices connect to a smart network, the city becomes increasingly exposed to vulnerabilities that can be incredibly difficult to resolve. While connected devices have the potential to help improve citizens' quality of life, sharing data could see the number of potential vulnerabilities that cybercriminals exploit increase. This is no surprise given our recent research which revealed that despite the majority of organisations having IoT devices on their corporate networks, more than 50% do not maintain necessary security measures beyond default passwords.

And so to mitigate this increased level of threat, towns and cities must ensure that they have a strong and reliable network. Without a secure and functional network, the technologies that make a city 'smart' would be redundant.

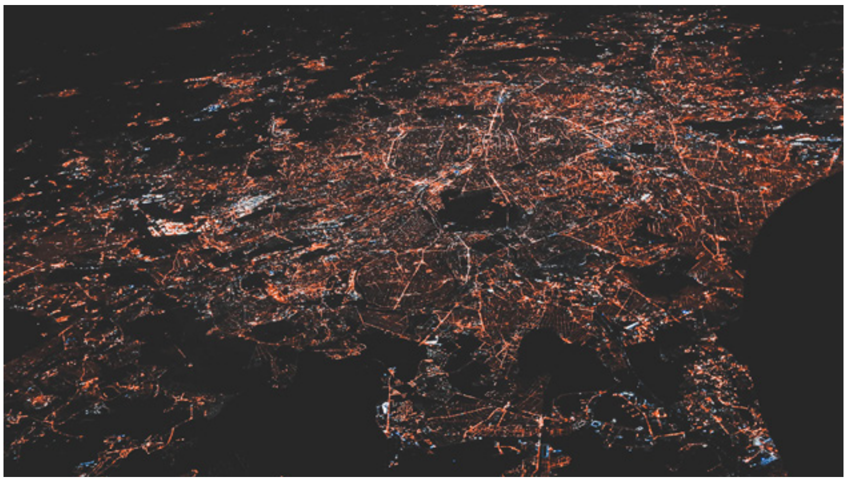

Countries around the world are looking to digitalise their infrastructures

A network ultimately provides the glue that sticks citizens and devices together. And so a fast, and reliable network, one that can transmit data efficiently, must be at the core of every smart city. After all, only by having an effective network will cutting-edge technology such as sensors, which can collect vast amounts of data to enhance the delivery of existing services to citizens and drive the creation of new services, be able to run.

But, it's not just about having a reliable network. The network itself must be secure, as if it were to be breached the use of devices would be hampered and worse, a whole city could grind to a halt. It's therefore essential that a secure network is the foundation of every smart city, to not only guarantee seamless usability but also safeguard against potential breaches.

To avoid costly errors, cities that are shifting towards becoming 'smart' must consider the importance of deploying a secure, robust, software-driven infrastructure. One of the easiest and most efficient ways is to establish multiple levels of access, thereby ensuring each user is segmented and contained. Segmentation of a network enables cities to create private virtual networks which allow critical services to be protected and isolated without any IP connection in or out. Taking such steps will both transform the user experience and ensure the network is secure and reliable.

Preventative measures can be put in place to ensure your city is not vulnerable and gaining full visibility to your network data is a key way to do this. Gaining full visibility, control and security over the IoT devices connected to the city's network will ensure that it can seamlessly function while remaining protected from potential cyberattack. ■

John Morrison, SVP of international markets, Extreme Networks

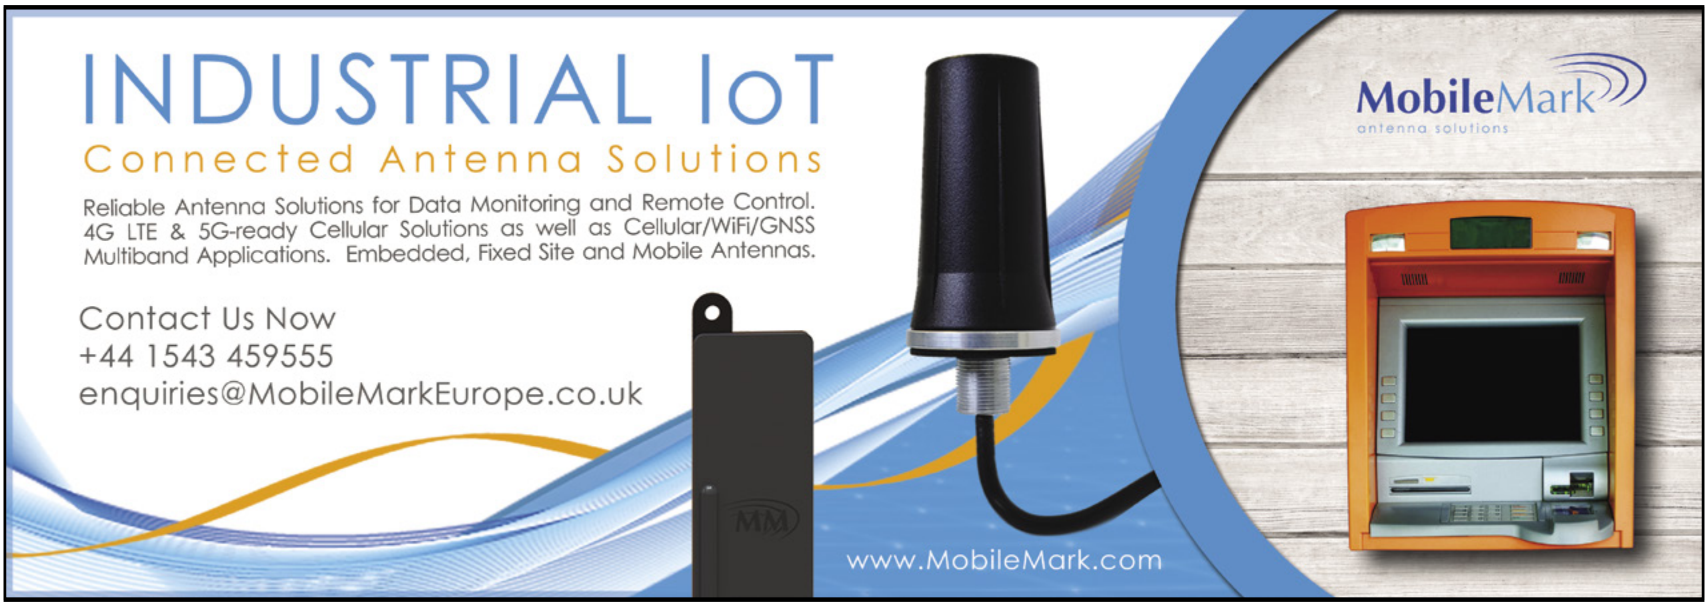

Data storage: bigger and faster comes at a cost

Brian Reed, vice president, product and alliances, Panasas

When describing the latest technologies, we use the phrase “state of the art” without being aware that we’re actually referring to a moving target.

Let’s consider high-performance computing (HPC) storage solutions, which help keep pace with the massive volumes of information that need processing.

Increasingly, HPC is central in tackling some of the most complicated tasks, from gene sequencing to vaccine development.

But describing a HPC system as state of the art doesn’t really account for all the factors that customers planning large-scale storage must have in mind. Not only are they expensive to buy, maintain and operate but the costs of downtime and outages are often overlooked until it is too late.

Users are waking up to this. While 57 per cent of HPC storage buyers said performance was the top criterion, 37pc mentioned long-term value or total cost of ownership (TCO) as a key factor according to a study by Hyperion Research for Panasas.

Familiar headaches HPC storage historically focussed on managing “big” files, whether a massive climate simulation or streaming files needed for CGI. Many relied on file systems that were ostensibly open source. Often platforms required more tuning after completing a job or preparing for the next.

But in the commercial world, there’s no tolerance for downtime and the staff required to keep things running.

Systems are expected to show a return on investment and handle multiple workloads simultaneously. In recent years, small files have played an increasing role, partly due to the demands of AI workloads, though anecdotally a similar pattern is being seen in traditional HPC areas such as life sciences and computational fluid dynamics.

Parallel file systems, with all the components talking to each other, were in danger of being swamped as the ratio of comms overhead to processing overhead increased.

The use of flash helped, but it is expensive compared to hard drives. One solution is to integrate flash and traditional storage. But that raises the challenge of managing the various tiers, to ensure they are used in the most performant way possible.

The cost of complexity, Hyperion highlighted, revealing insights about the total cost of ownership (TCO) as it applies to HPC.

One major cost is people. In 43 per cent of installations, one to three people were required to maintain it, while eight per cent needed four or five. Five or more were required at 10pc of HPC storage installations.

So, although just over a quarter of installations spent $100,000 or less in staffing, almost a third saw costs of $100,000 to $300,000, and almost 14 pc cost over $500,000. Simply recruiting and training experts was the most challenging aspect of HPC storage for 38pc of organisations.

Installation is also a major challenge. Just six per cent of organisations had their HPC storage rigs operational within a day, with 38% needing two to three. Over a quarter needed four to five days, with a similar number still unboxing after a week.

Downtime is also a major headache: almost half said they had to tune and retune systems monthly, with four per cent retuning weekly and two per cent daily. Additionally, monthly failures were reported by one-third of organisations and eight per cent reported weekly outages.

While 59% said recovery usually took a day or less, 24% took two-three days, while 14% took up to a week and three per cent needed more.

This is expensive, particularly for commercial customers adopting HPC storage, with 41% costing a day’s outage at up to $99,000. Fourteen percent put the cost at $100,000 to $500,000, with six per cent hitting $500,000 to $1M.

For four per cent, the daily outage cost was a shocking one million dollars.

With HPC storage installations expected to facilitate a wider variety of jobs, involving different file types, and with organisations developing a lower tolerance for failure, buyers will inevitably pay more attention to TCO.

Measuring the performance of HPC storage is an inexact science. There is a range of well-established parallel file systems, on a variety of hardware. Each installation is built for the specific needs of the client and its chosen applications

These are important considerations as HPC storage installations increasingly tackle a broader range of problems. The shortcomings of traditional approaches are becoming increasingly clear and it is harder to disguise or ignore the hidden costs of staffing and outages.

PRODUCTS

Just introduced by **Infinidat**, *InfiniBox SSA* is the company’s first to use 100 per cent solid state technology. The company claims that it delivers the industry’s highest performance for the most demanding enterprise applications.
It uses Infinidat’s algorithms and DRAM cache and is said to deliver performance and latency results that surpass all-flash arrays while providing the same customer experience, 100 per cent availability, and reliability of earlier models.
Infinidat says the new model, for those requiring ultra-low latency, complements the standard InfiniBox for general purpose applications needed by most enterprises.
And the common software of the two makes data mobility between them seamless, reducing the total cost of ownership. Both use the company’s cloud-based analytic tool, Infiniverse, to monitor and report on performance.
InfiniBox SSA has a claimed 546TB of usable capacity and an effective capacity of 1,092TB. Weighing 710g, it has a form factor of 26U and fits a standard 42U rack. *infinidat.com*

New hard drives from **Toshiba** are the company’s first to feature energy-assisted magnetic recording. This, it says, raises the capacity of each of the nine disks to 2TB for a total of 18TB – 12.5 per cent more than the previous models.
Helium sealed, the *MG09* range features what Toshiba describes as Flux Control - Microwave Assisted Magnetic Recording (FC-MAMR).
They are said to be compatible with the widest range of applications and operating systems and are adapted to deal with mixed random and sequential read/write workloads.
Toshiba says its new drives offer lower capital costs – leading to reduced total cost of ownership – for use in on-premises rack-scale operations and cloud and hybrid-cloud use.
With a form factor of 3.5in, able to fit standard bays, they are said to deliver 7200rpm, with a 550TB per year workload rating, and a choice of SATA or SAS interfaces. *toshiba.semicon-storage.com*

In one product, **StorageCraft** promises data storage which is easy to use, cuts costs and ensures reliable recovery.
OneXafe, it says, simplifies operations and eliminates management of disparate infrastructure by integrating the entire storage management and data protection stack into a single data infrastructure.
It is said to guarantee recovery with automated reverification of backup images and inflight verification of data.
StorageCraft says OneXafe instantly recovers backup images as virtual machines (VMs) in milliseconds with VirtualBoot I/O read-ahead technology.
And it says that OneXafe has the ability to recover to dissimilar hardware or virtual environments, ensuring recovery is timely while not waiting for specific resources.
OneXafe, says the company, is designed to offer scalable capacity for either primary or secondary workloads. It expands storage seamlessly, add one drive at a time, or multiple nodes within a cluster, without any configuration changes.
It says that OneXafe minimises storage and operational expenses with the use of powerful data reduction technologies such as inline deduplication and compression.

And it integrates with DRaaS to promise total continuity with a complete, orchestrated virtual failover to the cloud in case of disaster when used with the company’s cloud services. *storagecraft.com*

Designed to be affordable, StorCentric has introduced an entry level model in its Violin QV series of all-flash NVMe storage devices.
It says that the QV1020 has the best price/performance ratio and offers fast response even during peak processing.
And it is said to increase application availability and resiliency and ensure protection from outages.
The new unit offers a maximum usable capacity of 116TB and maximum effective capacity of 464TB. The earlier QV202 has a claimed maximum usable capacity of 479TB and 1.7PB of maximum effective capacity.
Both models, says StorCentric, owe their performance to Flash Fabric Architecture (FFA) that gives them high bandwidth and IOPs at consistent low latencies.
It says they offer 99.9999 per cent uptime; data protection with snapshots to enable complete recovery; deduplication for data reduction; web-based management; and data allocation, which writes to the NVMe SSDs are balanced across all drives for better performance. *storcentric.com*

Data generation is expected to grow at an annual rate of 42.2 per cent over the next two years, says **Seagate**, citing a survey it commissioned from IDC.

The company says its *Exos X18* drive helps to efficiently and cost-effectively manage ever-increasing amounts of data.

It describes Exos X18 as the fifth-generation high capacity helium 3.5in 7200 RPM nearline enterprise drive offering 18TB with both SATA and SAS interfaces. Seagate says it overcomes NAS system storage challenges by enabling ultra-fast data transfers, lower power and weight compared with traditional nearline drives and increased random reads/write (IOPS) than previous generations with 512e and 4KN formatting. And it offers built-in data protection through its Seagate Secure technology.

In addition, Seagate has introduced the new Exos Application Platform (AP) 2U12 as well as a new controller for the company’s AP 4U100 systems. *seagate.com*

Said to be 100 times faster than flash-based devices, SiliconDisk is a RAM-based 1U data storage unit from **ATTO Technology**.
It says *SiliconDisk* far exceeds current SSD products for performance and extensibility with under 600 nanoseconds of latency, four 100GB Ethernet ports and 25GB/s of sustained throughput.
And ATTO says it is plug and play, needing no special software or application or infrastructure changes.
Data is instantly stored and retrieved, says the company, making it ideal for accelerating real-time analytics. Memory capacity starts at 512 GB.
The four 100GbE channels are integrated into a single chip and linked to high-speed RAM, managed by the company’s XCOR storage controller designed to remove bottlenecks in performance.
Included with the device are RToptimizer, for real-time analytics, and Infinite Write Endurance, said to ensure that the RAM has no “per write” flash performance penalties or worry of memory wear-out. *atto.com*

“Please meet...

Gemma Moore, director and founder of Cyberis

What was your big career break?

Getting into cyber security was a bit of a revelation, but that happened very early on in my career! I did my degree in Computing at Imperial College, and even by the time I had graduated, I really wasn’t sure where to go with it. Development, architecture, systems administration, networking - I thought about lots of possibilities but none of them felt quite right for me. Specialism in cyber security wasn’t on my radar until I saw a job advertisement for a trainee penetration tester. It sounded like great fun, so I applied and was lucky enough to get the job. I’ve never looked back - I absolutely love working in this field, and penetration testing gave me exposure across all the disciplines I’d been interested in. It’s fast-paced, ever changing, and I never get bored.

Looking back, I don’t think there was a single ‘big break’ which pushed my career forward once I was in the field, but what I have benefitted from throughout is the supportive and collaborative atmosphere in the teams I’ve worked in.

Who was your hero when you were growing up?

This is a tricky one to answer, but I think it would have to be Terry Pratchett. I was a bookworm growing up, and still am, but there was never an author who felt more like a friend to me than Terry Pratchett. I devoured his Discworld novels, and they were full of his razor-sharp wit, empathy and wisdom.

His books tackled thorny subjects like structural injustices, corruption, prejudice, racism and superstition with a warmth and humanity that I’ve rarely felt from another author. There was so much of himself in his novels that jumping into a Discworld book felt like catching up with a friend - and still does.

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

My nan was full of good advice, but her favourite piece of advice for everybody was “JFDI” - Just F*****g Do It. “JFDI” gets you quite a long way when you’re the type of person who can spend too long thinking about things rather than doing them. Procrastinating? JFDI. Imposter syndrome? JFDI. Fear of failure? JFDI.

Lots of people, women particularly, worry too much about how others perceive them, or what might happen if they make a mistake, or how they will be judged and that worry can paralyse you. Sometimes, you need to acknowledge your feelings, shove them out of the way and just get things done anyway.

What’s the strangest question you’ve been asked?

Probably the one about whether I’d rather fight one horse-sized duck or a hundred duck-sized horses. I’d take the horse-sized duck every time.

What would you do with £1m?

I’m not sure I’d know until it happened! As a lump sum, that type of money could be transformative. I think living mortgage-free would be quite liberating, money doesn’t make you happy, but it does give you freedom and choices and using those wisely can make you very happy indeed.

If you could live anywhere, where would you choose?

I love living in the UK, but I’m very much a country mouse. I like peace and quiet, birdsong, open skies and the ability to walk my dog for miles from my front door. Preferably with super-fast broadband! I currently live surrounded by woodland which suits me perfectly. If I were to be persuaded to move from where I am now, it would probably be towards the Lake District or to the coastlines of Cornwall or Pembrokeshire. There’s such a variety of landscapes, I’d find it difficult to pick a favourite, so if money were no object, perhaps I’d have boltholes all over the country!

The Beatles or the Rolling Stones?

The Beatles.

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

I absolutely adore dogs, possibly in training or in behavioural consultancy is what I would want to do.

What’s the one thing you must do before it’s too late?

There’s a lot of travelling I would like to do that I’ve never managed to fit into life. I’ve never seen a rainforest for example, and it’s something I really want to experience at some point. The biodiversity, the wildlife, the bugs, the noise, the smells, the rain! ”