Inside_Networks

THE NETWORK INFRASTRUCTURE E-M

Nobody's perfect

WHY H
DATA C



DEC
26

UMAN ERROR IS A FACT OF LIFE AND HOW
CENTRE OPERATORS CAN MITIGATE IT



Beyond the back-up

HOW AI IS RESHAPING THE ROLE OF UPS SYSTEMS

A connection through time

HOW PRE-TERMINATED CABLING IS EVOLVING TO MEET THE DEMANDS OF HIGHER DENSITIES AND FASTER DEPLOYMENT



SIEMON™

BUSINESS CRITICAL AI IN FINANCIAL SERVICES

Tuesday 1st December | 10am GMT | 11am CET

As AI becomes increasingly embedded in financial services, network infrastructure must evolve to meet growing demands for performance, scalability and resilience.

Join Siemon experts Dominic Ross and Ryan Harris for this tech talk webinar as they explore the challenges of supporting distributed, regulated and real-time AI workloads, sharing practical insights from their experience working with financial service organisations.

Register Now

6 **ROB'S BLOG**
Present and correct



8 **NEWS**
All that's happening in the world of enterprise and data centre network infrastructures

12 **MAILBOX**
The pick of the recent emails to Inside_Networks

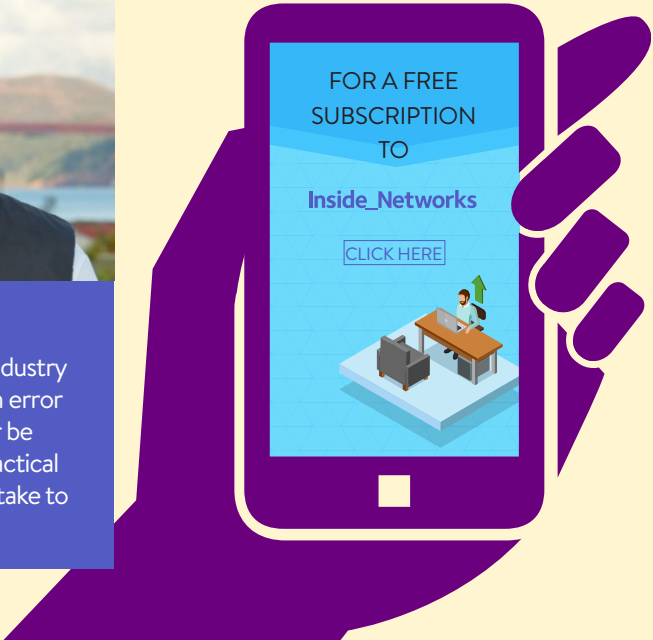


15 **QUESTION TIME**
Inside_Networks asks industry experts whether human error in data centres can ever be eliminated and what practical steps operators should take to mitigate it

24 **CHANNEL UPDATE**
Moves, adds and changes in the channel

30 **PRE-TERMINATED SYSTEMS**
In the first of a two part article on the past, present and future of pre-terminated cabling, Jason James of HellermannTyton looks at why it became essential

34 **PRE-TERMINATED SYSTEMS**
Christophe Hinet of Molex Connected Enterprise Solutions highlights the benefits of adopting a hybrid approach to pre-terminated and field terminated solutions



38 **PRE-TERMINATED SOLUTIONS**
A selection of the very best pre-terminated optical fibre and copper solutions currently available



40 **PRE-TERMINATED SYSTEMS**
Michael Akinla of Panduit identifies the key considerations when specifying and installing optical fibre and copper pre-terminated cabling

44 **KNOWLEDGE BANK**
Carrie Goetz examines optical fibre configurations for AI

46 **QUICK CLICKS**
Your one click guide to the very best industry blogs, white papers, podcasts, webinars and videos

48 **UPS AND POWER MANAGEMENT**
Ben Pritchard of AVK explains why the future of data centre power is modular



52 **UPS AND POWER MANAGEMENT**
Pierre Ecuyer of Legrand explains why AI is redefining UPS resilience

56 **UPS AND POWER MANAGEMENT PRODUCTS AND SYSTEMS**
Cutting edge UPS and power management products and systems profiled

62 **PROJECTS AND CONTRACTS**
Case studies and contract wins from around the globe

64 **PRODUCTS AND SERVICES**
The latest network infrastructure products, systems and services

67 **FINAL WORD**
Oskar Lampe of BCS Consultancy explains why South Africa is becoming a key hub for digital infrastructure and what this means for the continent's growing data centre market

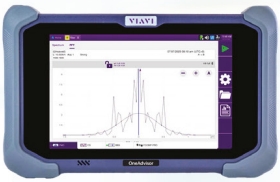
Inside the rack.
Across the campus.
Between data centers.

VIAVI Solutions

This is Fiber Testing at Hyper Scale

From high-density racks to long-haul interconnects, VIAVI brings simplicity to fiber testing for AI infrastructure at every stage.









viavisolutions.com/testatscale

To err is human

EDITOR
Rob Shepherd
07708 972170

SUB-EDITOR
Chris Marsland

ADVERTISING MANAGER
Kate Paxton
01603 610265

CREATIVE DIRECTOR
Vishnu Joory

TECHNOLOGY CONSULTANT
James Abrahams

CIRCULATION MANAGER
Debbie King

ACCOUNTS
Billy Gallop

All rights reserved.
No part of this publication may be used, transmitted or produced in any form without the written permission of the copyright owner. Applications for written permission should be addressed to info@chalkhillmedia.com
The views and comments expressed by contributors to this publication are not necessarily shared by the publisher. Every effort is made to ensure the accuracy of published information.
© 2026 Chalk Hill Media

Human error remains an unavoidable consideration in the operation of data centres, where even a seemingly minor mistake can have significant consequences. Uptime Institute's Annual Outage Analysis 2026 report identifies a failure to follow established procedures as the leading reason for human error related outages, raising important questions about how operators can better manage risk.

For this month's Question Time, Inside_Networks has brought together a panel of industry experts to explore whether human error can ever truly be eliminated. They examine the practical steps operators can take to strengthen procedures, improve working practices and prevent avoidable mistakes from developing into costly downtime.

Meanwhile, as data centres evolve, so too must the technologies and approaches used to support growing demands. In this issue, we turn our attention to pre-terminated cabling and Jason James of HellermannTyton looks at how pre-terminated systems became an essential part of modern cabling infrastructure, while Christophe Hinet of Molex Connected Enterprise Solutions considers the benefits of combining pre-terminated and field-terminated solutions. Michael Akinla of Panduit completes the line-up by examining the key considerations when specifying and installing both optical fibre and copper systems.

Power infrastructure is another area undergoing significant change, particularly as AI drives demand for greater capacity and resilience. Our special feature on UPS and power management brings together three industry perspectives on the challenges and opportunities ahead. Ben Pritchard of AVK explores why modular solutions are becoming increasingly important to the future of AI power and Pierre Ecuyer of Legrand then examines how AI is redefining expectations around UPS resilience. Last but certainly not least, Colin Dean of Socomec goes on to consider how operators can balance resilience with energy efficiency.

I hope you enjoy this issue of Inside_Networks and, as always, if you would like to comment on any of these subjects, or anything else relating to data centre and enterprise network infrastructure, I would be delighted to hear from you.

Rob Shepherd
Editor

Schneider Electric

Advancing Energy Tech

AI demands. **We Deliver.**

A complete portfolio built for the AI factory era.

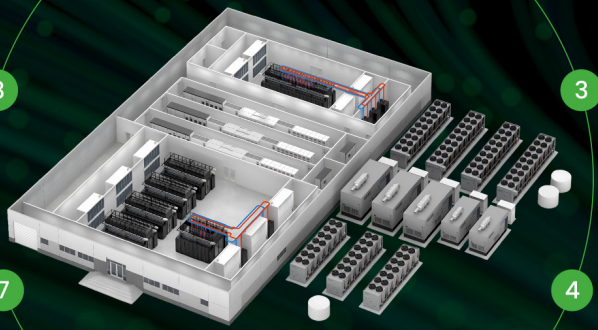
From grid interconnection to GPU cooling

From design simulation to AI-powered operations

From first deployment to global scale

Schneider Electric is the only partner with the portfolio, expertise, and ecosystem to power the full lifecycle of AI Factories.

1. AI Factory Reference Architecture & Co-Engineering
2. Digital Twin, Power & Thermal Simulation
3. End-to-end, grid-to-chip Power Products
4. Integrated Air & Liquid Cooling Infrastructure at Scale
5. High-Density Rack Architectures
6. Prefabricated & Modular AI Factory Deployment
7. Unified Controls & Automation
8. AI-Ready Software, Monitoring & Visibility
9. Lifecycle Services, Monitoring & Optimisation



Scan Me to Learn More

Nearly half of UK organisations cannot fully evidence where sensitive data is held and accessed

46 per cent of UK organisations cannot fully evidence either where sensitive data is backed-up or who can access it, research from Telehouse has revealed. The findings show many of the country's organisations cannot verify exactly where workloads are, or which suppliers or subcontractors have access.


Mark Pestridge

69 per cent have not yet reached the more advanced stages of preparedness for evolving data sovereignty, data protection and infrastructure regulation. In total, 76 per cent of organisations say they host at least some UK regulated data offshore. The research also highlights a gap in organisations' ability to demonstrate compliance across complex infrastructure environments. 34 per cent report inconsistent audit trails, supplier reporting or contractual assurances, while 28 per cent have only a limited view of suppliers and infrastructure partners.

Mark Pestridge, managing director at Telehouse Europe, said, 'Organisations are now managing data across complex networks of cloud platforms, data centres, service providers and back-up environments. Without clear visibility across that ecosystem, it becomes much harder to demonstrate compliance, provide customer assurance and respond effectively to changing regulatory requirements.'

Women and non-binary staff hold just 21 per cent of senior UK tech roles

Women and non-binary staff make up 29 per cent of UK tech employees but only 21 per cent of senior tech role holders, according to government data. The figures have been published as part of the work of the Women in Tech Taskforce, which works to address systemic barriers that prevent women from entering and leading in the tech sector.


Sheila Flavell

As the UK aims to scale the national AI workforce twentyfold by 2030, the sector already faces a shortage of 98,000-120,000 professionals across AI, cybersecurity and digital infrastructure. Yet rather than strengthening its talent base, the industry is steadily losing skilled professionals at a time of national urgency.

Sheila Flavell, chief operating officer at FDM Group, commented, 'Technology is changing the future of work, and if women aren't properly represented in how it's built and used, we risk hardwiring today's inequalities into tomorrow's systems. Organisations must create clear pathways for women into digital and AI careers, through structured training, upskilling and reskilling, and real opportunities to progress.'

Wireless Broadband Alliance sets out requirements to shape Wi-Fi and 6G convergence



The Wireless Broadband Alliance (WBA) has outlined requirements to help deliver more seamless connectivity across Wi-Fi and future cellular networks. According to the WBA, Wi-Fi is essential to delivering the 6G vision and should be considered in the design of the connectivity framework from the beginning.

Achieving the IMT-2030 ambition for ubiquitous connectivity means that a consistent user experience and deep indoor coverage will require Wi-Fi and cellular networks to work together more effectively. An objective is to create a more flexible, modular and access agnostic model that reflects how networks are deployed and operated, while making better use of the complementary strengths of Wi-Fi and cellular connectivity.

Tiago Rodrigues, WMA president and CEO, said, '6G must deliver ubiquitous connectivity across technologies, not simply a new cellular generation, and Wi-Fi will be fundamental to achieving that, particularly indoors. These requirements turn WBA's vision into a practical framework that can reduce cost and complexity, improve customer experience and coverage, and enable new commercial opportunities.'

Kao Data teams up with techUK and the DCA to launch Data Centres Explained

Kao Data has launched Data Centres Explained – datacentresexplained.com – a free educational resource providing information about data centres. Supported by techUK and the Data Centre Alliance (DCA), it clearly outlines what data centres are, why we need them and how they work, in clear and easy to understand language.



Visitors can explore a data centre zone by zone and at the heart of the platform is a public question portal, where anyone can submit a question through the website, with answers added so the resource grows in line with what people genuinely want to know. All content is informed by official data and guidance, industry standards and technical experts, is checked before publication and reviewed as technologies and policy continue to develop.

Matt Evans, chief operating officer at techUK, said, 'Data centres have only recently moved into the public eyeline and like most new infrastructure assets, they are often misunderstood. Data Centres Explained is exactly the type of open and clear engagement that help us all understand the role that they play in our lives.'

CBRE expands critical infrastructure capabilities in Europe and Asia Pacific with Onnec acquisition



CBRE Group has acquired Onnec Group. The acquisition will significantly enhance CBRE's ability to support its clients' data centre needs on a global basis. The initial purchase price is £460mn in cash and Onnec will operate as part of CBRE's Critical Infrastructure Services line of business. CBRE expects this business line to generate nearly $4.5bn of total revenue in 2027 including contributions from Onnec.

Founded in 1986 and based in London, Onnec specialises in the design, deployment and maintenance of network, optical fibre and data centre infrastructure. Approximately 800 employees serve hyperscalers across key markets and Onnec's geographic presence aligns with the data centre infrastructure services that CBRE delivers in the US.

'CBRE is pursuing an explicit strategy to build a global infrastructure services business to complement our real estate services business,' said Vikram Kohli, CBRE's chief operating officer. 'We've long believed Onnec is an excellent fit in this strategy and the acquisition is also consistent with our ongoing effort to build resilient businesses that benefit from secular tailwinds.'

CommonAI selects Kao Data to host the UK's Scaling Inference Lab

Kao Data has been chosen by CommonAI to host its pioneering AI Scaling Inference Lab at its data centre campus in Harlow. Funded by the Advanced Research and Invention Agency (ARIA) and delivered by CommonAI, the lab provides a neutral environment where emerging AI technologies are put to the test and benchmarked against commercial systems.

CommonAI will deploy a series of high density, graphics processing unit (GPU) accelerated AI clusters within two dedicated technology suites. Each cluster tests a distinct approach to AI infrastructure including pipeline parallelism, heterogeneous orchestration and nanosecond optical circuit switching, with all five clusters aiming to become operational by early 2027.



'We are proud to become home to the Scaling Inference Lab,' said Spencer Lamb, chief executive officer at Kao Data. 'The work taking place at our Harlow campus will help shape the next generation of UK AI, ensuring the country remains a leader in AI development while meeting the growing demand for sovereign, homegrown AI infrastructure and the resources needed to support the country's world leading start-up community.'

82 per cent of UK and Ireland finance leaders feel under pressure to adopt AI faster than they are comfortable with

82 per cent of finance leaders in the UK and Ireland say they feel under pressure to adopt AI in their finance function faster than they are currently comfortable with, according to research from AccountsIQ. The CFO Mindset 2026 research, conducted by Censuswide among 1,036 finance leaders in the UK and Ireland, found that 41 per cent describe the pressure as significant.



The pressure is highest in Ireland, where 91.5 per cent of finance leaders feel it, against 80.3 per cent in the UK. The pressure sits alongside high trust in the technology, with 81 per cent of finance leaders saying they fully or mostly trust AI-driven outputs in finance processes.

Darren Cran, CEO of AccountsIQ, said, 'Finance leaders are being asked to move on AI at the speed of the conversation around it. The discussion of AI starts in the boardroom and the press, and finance is expected to have an answer. Most finance leaders already trust what AI produces, so the hesitation sits with how quickly they can bring it in with the controls they need.'

NEWS IN BRIEF

Virtus Data Centres has completed a £2.45bn financing package. The capital provides a long-term funding framework to support the continued development and expansion of Virtus' data centres across the UK and Europe and includes a £1.2bn green capital expenditure facility, available through both term and revolving tranches, providing flexibility to fund expansion and support its growth ambitions.

Research by global law firm White & Case reveals that 97 per cent of UK defence companies and defence technology start-ups surveyed are likely to raise significant external funding over the next 12 months to accelerate growth and meet soaring demand for defence technologies.

As more travellers turn to AI tools to research destinations, build itineraries and plan holidays, research suggests travellers are still reluctant to hand over complete control. Booking.com found that 67 per cent of travellers have already used AI for some part of a trip, while 89 per cent say they want to use the technology to help plan future travel. However, just six per cent say they fully trust AI, and only 12 per cent are comfortable with the tool making decisions independently.

Fibre Council Europe has unveiled Fibre Horizons 2027, which takes place on 17th-18th March 2027 at Allianz MiCo in Milan. Fibre Horizons builds on the legacy of the FTTH Conference, carrying forward more than 20 years of industry recognition while evolving the event to reflect the changing role of optical fibre within Europe's digital infrastructure ecosystem.

Why Europe's fibre boom dep

Hi Rob



Europe's optical fibre industry has become exceptionally good at counting homes passed. The figure is easy to understand, useful for comparing markets and politically attractive. However, it is no longer sufficient evidence of success.

According to Fibre Council Europe and IDATE, fibre coverage across the EU39 has reached 79.3 per cent, while take-up stands at 54.4 per cent and household penetration at 43.1 per cent. These figures reveal the challenge now facing the sector – a substantial proportion of the premises that can theoretically obtain fibre have not become active connections.

This matters because fibre networks were not financed simply to pass buildings. Governments supported deployment to advance digital inclusion and economic development. Commercial investors funded networks in expectation of long-term subscription revenue. Neither objective is fully achieved when fibre reaches the street but stops short of the customer.

The industry therefore needs to reconsider what it measures and rewards. Homes connected, time to activation, cost per installation, service quality and customer retention now tell us at least as much as headline coverage.

The remaining barrier is often described as the 'last 100m' but distance is only part of the problem. The final connection can involve building owners, landlords, housing associations, tenants, municipalities and installation partners. In apartment buildings, technicians must negotiate building entry points, risers, floor distribution and individual dwellings. Permissions, incompatible infrastructure, inaccessible spaces and missed appointments can delay an activation more effectively than kilometres of civil engineering.

These complications help explain why strong rollout does not automatically produce equally strong adoption. The UK, for example, reported full-fibre availability at 78 per cent of residential premises in July 2025 but take-up among premises with access was 42 per cent. Germany and the UK have both continued to add millions of homes passed, yet the commercial focus must increasingly turn from construction volume to conversion.

ends on the last 100m



That requires a different operational mindset. The best network is not necessarily the one that was cheapest to build past the greatest number of properties. It may be the one that can connect each new customer quickly, predictably and with minimal disruption.

Standardised building designs, modular distribution points, pre-terminated connections, accurate documentation and simpler field procedures can all reduce activation time and dependence on scarce specialist labour. The objective should be to avoid unnecessary site visits, reduce errors and allow installers to complete more connections without compromising quality.

The customer's experience must also be considered. Consumers do not purchase fibre as a civil engineering achievement. They purchase dependable video calls, streaming, cloud access, gaming, telemedicine, online education and smart home services. Their judgement of the network is formed inside the home. Poor internal connectivity or unreliable Wi-Fi will be attributed to the broadband provider – however impressive the outside plant may be.

Rural deployment presents an additional test. Public programmes are rightly extending fibre into areas that commercial investment might not otherwise reach. However, long routes, dispersed properties and lower customer density make efficient activation especially important. Reusing existing poles and separating network deployment from the connection of individual subscribers can help control costs but coverage targets should still be accompanied by credible plans for adoption.

Europe's first fibre challenge was getting infrastructure into the ground. Its next is turning that infrastructure into services people use. That's why policymakers, investors and operators should place far greater emphasis on the journey from availability to activation.

A home passed represents potential. A home connected represents a working service, a customer relationship and a return on investment. That is the distinction that should shape the next phase of Europe's fibre strategy.

Thomas Ritz
R&M

Editor's comment

Thomas' comment that 'consumers do not purchase fibre as a civil engineering achievement' neatly captures why ambitious coverage targets must be matched by credible, customer focused adoption plans. Fibre success isn't measured by how many homes are passed but by how effectively those homes are turned into connected customers.



The leading trade body for electrotechnical contractors

STRENGTHEN YOUR BUSINESS

THROUGH ECA MEMBERSHIP

- Recognised professional credibility across public and private sector clients
- Technical and business-critical expertise to support compliance, reduce risk and protect margins
- Tailored savings and services that deliver measurable value
- Access to a powerful peer network, shaping industry standards and policy

For electrotechnical contractors committed to excellence, resilience, and growth, ECA membership is essential.

Discover more at eca.co.uk/unlock

ECA

Make no mistake

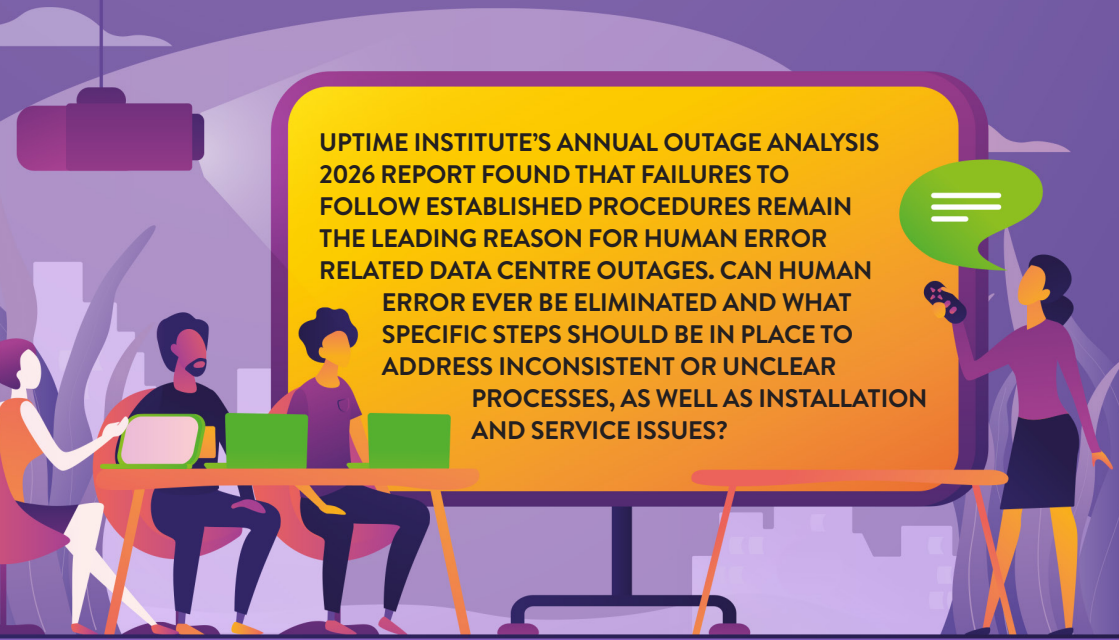
The figures quoted for data centre outages caused by human error have been high for several years. Inside_Networks has brought together a panel of industry experts to examine whether human error can ever be eliminated and what practical steps operators should take to mitigate it

Human error remains an unavoidable consideration in data centre operations but the consequences of mistakes can be anything but minor. While technology, automation and monitoring can help reduce operational risk, people remain responsible for designing, installing, maintaining and managing critical infrastructure.

Uptime Institute's Annual Outage Analysis 2026 report highlights failures to follow established procedures as the leading reason for human error related outages. This raises questions about whether the issue lies with individuals or whether unclear, inconsistent or difficult to follow processes are creating the conditions for mistakes to occur.

Beyond operational procedures, installation and service activities can introduce further risks, particularly where standards, documentation, training or accountability are inconsistent. As data centre environments become increasingly complex, ensuring that processes are clear, repeatable and properly followed is becoming ever more important.

Inside_Networks has assembled a panel of experts to discuss whether human error can ever truly be eliminated. They explore the practical measures operators can put in place to reduce risk, improve processes and address issues before they result in downtime.



MATT RENNER

CHIEF TECHNOLOGY OFFICER AT SALUTE

There will always be error when there are humans involved. That is obvious and eliminating it is the wrong goal. Our focus instead should be to build systems that assume it will happen and maximise our ability to mitigate it.

What matters is that our industry's operational leadership puts our peers in a safe environment and in a position to be successful. Uptime's Annual Outage Analysis 2026 report names failure to follow established procedures as the main culprit, which raises two concerns:



- Are operators consuming procedures in a way they can safely execute?

  Recalling a procedure from memory months after onboarding is inherently risky. This is why we invest heavily in our workforce through structured training and partnerships with specialist providers to drive meaningful career progression. This approach is foundational to the people infrastructure our teams need to be successful.

- Are the procedures themselves engineered correctly?

  In our industry, a procedure is often written once at handover and inherited unchanged through operator transitions. That is how unclear process takes root and it is usually not discovered until someone is in front of the equipment trying to execute it, or worse, after an outage. If something goes wrong, it gets logged as human error, even though the procedure was the root cause.

Together, these two concerns form the feedback loop of operational intelligence. Every task our teams execute tells us something – where a step caused concern, where a technician stopped to ask a question or where a near miss almost became an event. We owe it to our peers and our customers to build the systems that capture those learnings and feed them back into the procedure and the training behind it.

That is the connection between what one of our people does today and how the next person does it better tomorrow. At our scale, that only works when the technology systems and the people infrastructure are tightly coupled.

Human error will not disappear. But where every execution makes the next one safer, that error becomes rarer, more visible and more recoverable.

> 'IN OUR INDUSTRY, A PROCEDURE IS OFTEN WRITTEN ONCE AT HANDOVER AND INHERITED UNCHANGED THROUGH OPERATOR TRANSITIONS. THAT IS HOW UNCLEAR PROCESS TAKES ROOT AND IT USUALLY IS NOT DISCOVERED UNTIL SOMEONE IS IN FRONT OF THE EQUIPMENT TRYING TO EXECUTE IT, OR WORSE, AFTER AN OUTAGE.'

MIKE MEYER
EXECUTIVE DIRECTOR AT THE DCA

People design, build, run and maintain critical infrastructure. Of course, mistakes happen. However, Uptime Institute's Annual Outage Analysis 2026 makes an awkward point – the main cause of outages linked to human error is still people not following procedures. That's not a technology problem and it is mostly within our control.



We've worked out how to engineer for 99.999 per cent availability. What we haven't cracked to the same degree is how people work. A facility can run at five nines, so why shouldn't the team behind it aim for six – 0.0001 per cent human error? Zero isn't realistic. Very close to zero might be, if the management team is steady enough.

I'd start with the obvious – make the process clear. If it is muddled or changes from one site to the next, then blaming the technician with the checklist gets us nowhere. There should only be one version controlled library of standard operating procedures (SOPs), method of procedures (MOPs) and emergency operating procedures (EOPs). The people using those procedures should help write them. Then test them in real conditions and go back to them after every incident and near miss. If someone must guess, the procedure has failed.

Training is the other part of it. Not the onboarding box tick but proper drills and failure scenarios. There needs to be a clear sign-off before anyone goes near a live system, followed by regular recertification. Do it often enough and the procedure stops being a document and starts becoming instinct. We should rehearse the day they never want to have.

Installation and servicing bring their own problems. Ideally, there would be one standard across the supply chain. Contractors and original equipment manufacturer (OEM) engineers need to work to the site's procedures, under clear permit to work controls, with a named site lead accountable for what happens. And commissioning? That isn't the paperwork at the end. Management should own it.

Culture is where all of this either works or falls apart. People need to be able to stop a job, challenge an instruction and say when they've got something wrong. No drama. No hunt for someone to blame. Leaders should reward honesty, develop their operations people and hire for judgement, not just technical skill. That's how a mistake gets caught before it becomes an outage.

We're not taking people out of the data centre and nor should we. The job is to lead them well enough that human error becomes the rarest nine of the lot.

> 'PEOPLE NEED TO BE ABLE TO STOP A JOB, CHALLENGE AN INSTRUCTION AND SAY WHEN THEY'VE GOT SOMETHING WRONG.'

ED BISSELL

SALES AND MARKETING DIRECTOR AT STELLIUM DATACENTERS

Sad but true, human error is typically the root cause of a whole range of potential unplanned outages, ranging from accidentally tripping the emergency power-off switch and unplugging live cables to overloading electrical circuits during maintenance. Furthermore, incorrect software or hardware configurations during routine upgrades or system changes can be an issue. Another is maintenance oversights where contractors or cleaning staff forget to isolate fire suppression/alarm systems, mishandle raised floor panels or leave physical access points unsecure.



Recent Uptime Institute research found that of reported major human error related outages, 58 per cent cited staff failing to follow procedures and 45 per cent cited incorrect processes or procedures. It proposes the most effective approach to reducing human error is by treating it as a system design problem, not just a training issue. I agree.

Continuous training should, of course, be mandatory for both technical personnel and third-party facility suppliers. But to further mitigate human error requires the enforcement of strict documented, site specific method of procedures before any engineer or third-party vendor accesses critical infrastructure. Investment in the upgrading of existing monitoring automation tools is another key consideration, making them more user friendly with displays that are clear to read, thereby reducing fatigue.

To err is human, so, yes, AI can help too – but only as a control layer around human operations, not as a replacement for experienced operators. For example, applying AI to generate realistic scenarios and testing whether operators can correctly execute critical procedures. Use it to monitor and compare procedural consistency, to identify conflicting instructions, missing steps, outdated equipment names or ambiguous language.

It could also help with procedure creation, turning approved engineering requirements into standardised, step-by-step procedures. Incident/near miss analysis is another task for AI, to help spot recurring procedural weaknesses by analysing logs, procedures, work orders and incident reports.

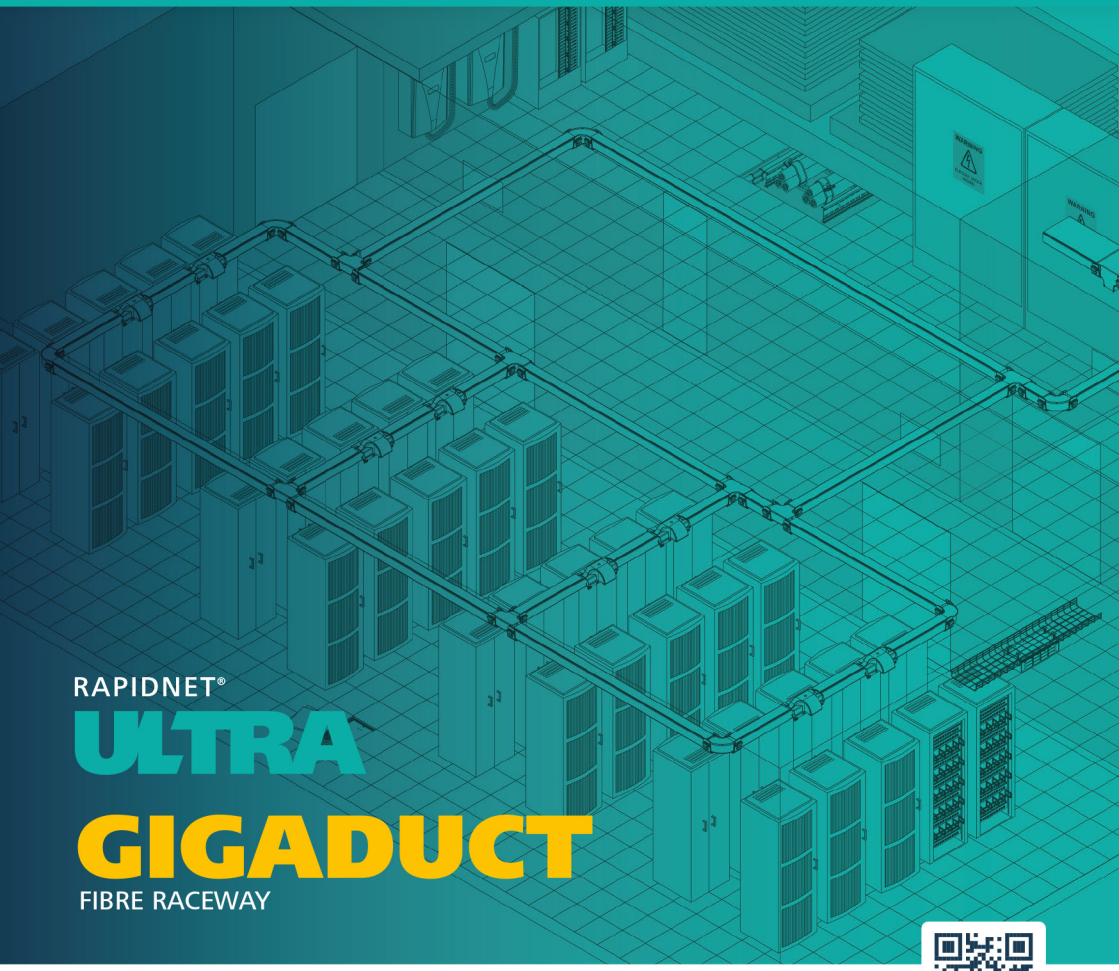
It could also be used to conduct 'what if' scenarios such as checking before a maintenance activity that the proposed method of procedure tallies with the equipment state, planned work, dependencies, previous changes and required approvals. Following maintenance, AI could also conduct post-job verification, ensuring the required steps, readings, alarms and redundancy states have been restored.

'TO ERR IS HUMAN, SO, YES, AI CAN HELP TOO – BUT ONLY AS A CONTROL LAYER AROUND HUMAN OPERATIONS, NOT AS A REPLACEMENT FOR EXPERIENCED OPERATORS.'

Next-gen Solutions for Fibre Connectivity

Integrating RapidNet ULTRA with GigaDuct unlocks a next-generation fibre management ecosystem: ultra-compact connectivity paired with a future-ready raceway system designed to protect, optimise and elevate data centre performance.

Click to find out more



HellermannTyton

www.htdata.co.uk/rnu

STEPHEN BOWES-PHIPPS

VICE PRESIDENT EMEA DATA CENTRES AND CLOUD AT STATE STREET

The industry has professionalised hugely during the last 15-20 years. Best practice is often shared among even the smaller boutique data centres as staff have often moved around from provider to provider, taking their operational nous with them. As plant and equipment have also become so much more stable, with equipment failures much rarer than in the 2000s, human error is going to play a much larger part in the availability of the data centre.



There are two types of human error. The first is a failure to follow procedure and the second is an inability to follow procedure.

The first is easily resolved. Hiring decisions should be based not only on competency but attitude as well. Don't put an individualist maverick, no matter how experienced and competent they may be, in a critical role where they are the last line of defence against failure.

The second is much harder to tackle and relies on experienced, passionate process engineers within the organisation. I wouldn't be exaggerating if I said I have yet to find, in all the organisations I have worked in, many technical people who understand the difference between a process and a procedure.

Complexity in documentation is a real cause of outages. Live by the mantra 'keep it simple stupid' (KISS). If a procedure is more than a page long, it's unlikely to be followed every time.

Procedures, in the form of standard operating procedures (SOPs), only work well if they are part of an established process. By established I mean written down with a documented agreed workflow. SOPs should also be trialled regularly and used for training (new staff and annual refresher). Don't hang your team out to dry and make sure they have exception processes for when the SOPs fail to cover specific circumstances. 'Making it up as they go along' is a recipe for disaster!

Technology can help and no doubt we'll see the rise of smart glasses that are able to access relevant documentation for engineers and maintenance staff in real time as they walk around the facility. But the basic principles remain – KISS and keep staff updated or you'll get failures.

'THERE ARE TWO TYPES OF HUMAN ERROR. THE FIRST IS A FAILURE TO FOLLOW PROCEDURE AND THE SECOND IS AN INABILITY TO FOLLOW PROCEDURE.'

RICHARD COLLAR

VICE PRESIDENT OF CUSTOMER SOLUTIONS AT KAO DATA



Human error cannot be eliminated from data centre operations and it is not helpful to pretend otherwise. People design, build, commission, maintain and operate these facilities – and will occasionally make mistakes. The objective is not to remove the human element but to create environments where errors are less likely, easier to detect and less able to escalate into service impacting incidents.

The Uptime Institute's finding that failure to follow established procedures remains a leading cause of human error is consistent with complex operational environments. Procedures are typically ignored, misunderstood or worked around when unclear, out of date, too generic or not aligned with actual plant and operating conditions. If a method statement or switching procedure does not reflect the asset in front of the engineer, it quickly loses authority.

This must be addressed from the earliest stages of design. Resilient infrastructure should assume that people may occasionally select the wrong device, follow the wrong sequence or act under pressure. Clear segregation, intuitive layouts, consistent labelling, defined isolation points and concurrently maintainable architectures reduce the probability that a single error leads to an outage.

Procedures also need to be owned by the people who use them. They should be written with operations teams, tested through practical walkthroughs and updated whenever the facility, control philosophy or maintenance regime changes. Version control, technical approval and periodic drills are essential if documentation is to remain current and trusted.

Installation and service activities require the same discipline. Robust commissioning, witnessed testing, accurate as-built information and clear handover packs establish the baseline for safe operation. Thereafter, permit to work controls, change management, risk assessments and approval gates provide structure around intrusive work – particularly where live systems are involved.

Culture is equally important. Engineers and service partners must feel empowered to challenge an instruction, stop a task or report a near miss without fear of blame. A mature operator treats these moments as leading indicators, not administrative noise. By combining resilient engineering, practical procedures, disciplined change control and a learning culture, the industry can reduce human error risk materially, even if it can never remove it entirely.

'PROCEDURES ARE TYPICALLY IGNORED, MISUNDERSTOOD OR WORKED AROUND WHEN UNCLEAR, OUT OF DATE, TOO GENERIC OR NOT ALIGNED WITH ACTUAL PLANT AND OPERATING CONDITIONS.'

BEN CRANHAM

CHIEF OPERATING OFFICER AT PULSANT

The latest Uptime Institute Annual Outage Analysis again highlights human error as a significant factor in service outages across the industry. But can we really blame the one thing that's guaranteed – people's fallibility?



We all mess up. It doesn't matter how experienced we are, how well trained or how long we've worked in critical environments. In fact, errors often come from opposite ends of the experience spectrum. The newest members of a team may not yet have the knowledge or context needed to make the best judgement. Equally, the most experienced can fall into the trap of familiarity, making assumptions or overlooking something because they've seen a similar situation countless times before. We're safest somewhere between competence and complacency.

Our responsibility as operators is to design, build and run services that benefit from the elegance of the human touch, while ensuring that the most human of things, an error, doesn't result in a service impact. Too often, discussions about human error focus on the individual. In reality, service impacts rarely come down to a single mistake. More often, a mistake simply exposes weaknesses that already existed.

Peer reviews, checklists, change advisory boards, competency assessments and regular retraining all play a role. Even highly capable people benefit from challenge and independent scrutiny. The best peer reviews aren't box ticking exercises. They create opportunities for someone to ask the question everyone else missed.

Technology can strengthen those controls through automation, validation, monitoring and approval workflows. But this isn't an argument for taking human skill out of the loop. The judgement, experience and expertise of operational teams remain fundamental to resilient service delivery.

Resilience can't be assumed just because a service was designed to be resilient. Changes accumulate, equipment ages, dependencies evolve and documentation drifts. Teams change and knowledge moves with them. Over time, the resilience that exists in practice can differ from the resilience that was calculated on paper. As operators, we need to continually test whether the resilience we started with is still going to be there when we need it.

The organisations that consistently deliver reliable services are not those that expect people to be perfect. They are the ones that combine skilled people, robust processes, constructive challenge and technology to ensure that humans remain the strongest part of the system, not its weakest link.

> 'IN REALITY, SERVICE IMPACTS RARELY COME DOWN TO A SINGLE MISTAKE. MORE OFTEN, A MISTAKE SIMPLY EXPOSES WEAKNESSES THAT ALREADY EXISTED.'

molex

The Field Lessons of 2026

12 Preventable Issues and How to Avoid Them

Connector contamination, unclear port mapping and crowded pathways can turn routine cabling work into extended troubleshooting. Knowing what to check first helps teams investigate the cause, protect adjacent live services and identify what needs to change to prevent a repeat.

The Field Lessons of 2026, from Molex Connected Enterprise Solutions, examines recurring issues in copper and fiber optic structured cabling, with practical guidance for installation, maintenance and handover.

Discover how to:

- Make connector inspection and careful handling part of everyday work.
- Verify labels, port mapping and end-to-end fiber optic polarity.
- Maintain cable support and access without disturbing neighboring connections.
- Match test results to installed links and resolve documentation gaps before handover.



Find the Cause. Prevent the Repeat.

Check the link before replacing hardware

Change one connection without disturbing another

Know exactly what passed testing

Download the Field Lessons Guide

866-733-6659

inquiries.ces.amers@molex.com

www.molexces.com

Bruce Claassen joins Kao Data as new chief financial officer

Bruce Claassen has joined Kao Data as its chief financial officer. He will help scale the business in support of its ambitious plans while partnering closely with the board and investors as new developments progress in Stockport and London, alongside the company's campus in Harlow.


Bruce Claassen

With more than 20 years of experience leading finance organisations across multinational businesses, Claassen brings a strong track record in capital investment, refinancing and building high performing finance teams across both listed and private equity backed businesses.

He joins from Global Switch, where he spent four years as group finance director and most recently acting chief financial officer.

'I am excited to be joining one of the UK's leading digital infrastructure businesses as it continues to support growing demand from the AI and hyperscale cloud sectors,' said Claassen. 'I look forward to working to deliver Kao Data's long-term growth strategy, building on the company's strong foundations and reputation for developing and operating data centres engineered for AI and advanced computing workloads.'

EfficiencyIT appoints modular infrastructure expert Graham Robson to scale prefabricated data centre delivery

EfficiencyIT has appointed Graham Robson as its new technical estimations manager. The appointment strengthens the company's engineering and operational leadership as AI and high density compute drive growing demand for prefabricated data centres that can be deployed quickly, efficiently and to specification.


Graham Robson

With more than four decades of engineering experience, including over 10 years dedicated to modular infrastructure, Robson brings deep expertise in modular design, critical power systems and mechanical and electrical (M&E) project delivery. His career has covered almost every stage of the modular build process, giving him a rare understanding of how prefabricated systems are designed, integrated and fitted out at scale.

Robson said, 'What drew me to EfficiencyIT was the calibre of its leadership. Nick, Ross, Graham and the team are innovative, firmly focused on tomorrow and measure success by the outcomes they deliver, taking every customer on a journey towards genuine efficiency and excellence. That is a rare commodity in today's market.'

BCS Consultancy launches new division to tackle data centre power constraints

BCS Consultancy has launched its new power infrastructure division to help investors, developers and operators overcome what has become one of the biggest barriers to data centre growth – securing a credible and deliverable route to power. The division reflects a fundamental change in both the scale and complexity of the power challenge data centre developments and projects are currently facing.


Cameron Thompson

While BCS' power activity was previously weighted towards delivering traditional substations for projects with established grid connections, client demand has increasingly moved upstream towards power strategy and identifying credible routes to unlock developments earlier. Led by Cameron Thompson, the team works with investors, landowners, developers and operators from early stage power strategy and grid connection through to the delivery of substations, cable routes, microgrids and other critical power infrastructure.

Thompson said, 'Power has always been fundamental to data centre development but the conversation has shifted. Clients are coming to us much earlier, asking how they can make sites viable within the timeframes they need.'

CHATSWORTH PRODUCTS

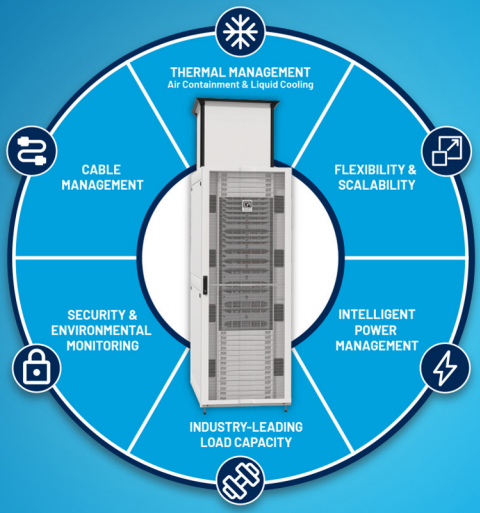
Discover More...

CPI STRATEGIC PARTNER
ZUTACORE

Optimise Your Data Center

Maximise Performance with an Integrated Cabinet Solution

Optimise space, power, cooling, and more with the ZetaFrame® Cabinet Systems fully integrated solution designed for high-density deployments, AI, and HPC.



Datalec strengthens its liquid cooling capabilities with the appointment of Dominik Dziarczykowski

Datalec Precision Installations has appointed Dominik Dziarczykowski as business development manager, strengthening the company's expertise in liquid cooling and high density data centre infrastructure.


Dominik Dziarczykowski

Dziarczykowski brings more than 20 years of experience across the IT, consulting and data centre industries, with a particular focus on high density computing, liquid cooling and data centre efficiency. His career includes experience with organisations including IBM, Dell, Huawei, Deloitte and Capgemini, as well as specialist roles in liquid cooling and data centre infrastructure. He previously worked with CoolIT Systems and most recently served as market development manager high density and liquid cooling at Vertiv.

Dziarczykowski commented, 'I'm excited to join Datalec at a time when AI and high density computing are reshaping data centre infrastructure. I look forward to working with our customers and partners to develop practical liquid cooling solutions.'

JBA appoints Justin Neil as director of technical and design

JB Associates (JBA) has appointed Justin Neil as director of technical and design. He takes up the new role with a remit to lead and build on the company's design capability, technical quality and project delivery whilst contributing to its strategic direction and future growth ambitions.


Justin Neil

Neil brings with him over 25 years of sector experience combining technical expertise, engineering leadership and commercial awareness. He has previously held senior roles within Hoare Lea, WSP and Deerns, and spent four years working in Qatar. He is a chartered engineer with extensive experience across the built environment and engineering sectors.

Commenting on his decision to join JBA, Neil said, 'I have known Ashley Buckland for more than 20 years and we worked together at NG Bailey early in our careers. He and the team have built a really strong business and have ambitious growth plans. I am excited to take on this role and help support the next phase of the company's development and strategic direction.'

CHANNEL UPDATE IN BRIEF

Opengear has signed its first pan-European distribution agreement with Ingram Micro. The agreement will give resellers in participating Ingram Micro countries access to Opengear's portfolio, supported by Ingram Micro's regional infrastructure, logistics capabilities and partner network. This expands Opengear's existing reach in the UK, Ireland, Spain, Germany, Italy, France, Belgium and Luxembourg, and brings its network resilience platform to resellers and end users in additional European markets.

Black & White Engineering has opened two offices in the US. Virginia and Florida form the first phase of the expansion, creating a local presence from which Black & White will support clients and projects across North America. The new locations take its global network to 34 locations with plans for an office opening in Texas already underway.

Microchip Technology and Navitas Semiconductor have collaborated on an 800VDC to 6VDC reference design for AI data centre rack power applications. The platform combines Microchip's digital power control and security technologies with Navitas' GaNFast gallium nitride (GaN) power devices to give developers a practical path to implement high efficiency power conversion aligned with the Open Compute Project (OCP) 800VDC standard.

Vertiv has expanded its Nové Mesto nad Váhom campus in Slovakia, reinforcing the company's commitment to supporting customer demand across Europe, the Middle East and Africa (EMEA). Michael Resha has also joined the company as executive vice president of manufacturing, logistics and operational excellence.



Best practice for validating data centre installations

As hollow core fibre (HCF) moves from lab research into field trials and early operationa manager for fibre optic field solutions at VIAVI, explains why getting validation right firs

HCF offers data centre operators a genuine step-change over traditional singlemode fibre (SMF). Unlike SMF, HCF guides light through a hollow channel surrounded by a ring of nested glass capillary tubes confining most of the transmitted light to within the air core. This substantially reduces – though does not fully eliminate – its interaction with glass.

See the light

The payoff is significant, as light travels faster and close to its velocity in a vacuum. This means up to 33 per cent latency reduction, up to 65 per cent lower attenuation for longer non-amplified links, and up to 70 per cent lower chromatic dispersion for fewer transmission errors, as well as reduced non-linear transmission effects. This all makes HCF an interesting prospect for financial trading organisations, defence applications, quantum cryptography and AI data centres.

But HCF cannot be validated using traditional SMF optical time domain reflectometer (OTDR) test approaches. Four physical effects complicate matters:

- Rayleigh backscatter (RBS) is 14-20dB weaker than SMF, reducing OTDR trace visibility and limiting event detection.
- The backscatter coefficient is non-uniform along the fibre, which complicates unidirectional trace interpretation.
- HCF-to-HCF splices introduce strong reflections that create OTDR dead zones.
- HCF's heightened bend sensitivity generates attenuation that is easily mistaken for poor splices or masked altogether by fluctuating backscatter.

Compounding this, field splicing briefly exposes the open core to ambient air, allowing gas absorption regions to form that further distort OTDR traces. Left unchecked, these effects risk false passes and problems surfacing only after deployment.

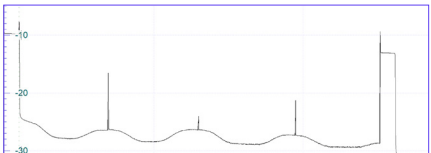


Figure 1: OTD trace showing non-uniform RBS coefficien and reflective fibre transitio splices in HC

Analyse this

The solution is bidirectional OTDR analysis. Combining measurements from both ends suppresses backscatter related artefacts to reveal the link's true loss profile. However, standard bidirectional algorithms alone are not sufficient, as reliable HCF characterisation, qualification and certification demands custom OTDR post-processing.

The same HCF physical properties also raise the bar for OTDR hardware. Low and variable RBS combined with highly reflective

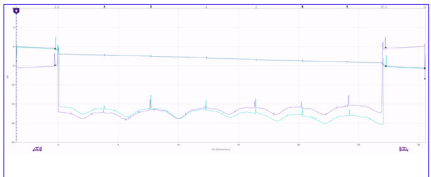


Figure 2: Bidirectional OTDR traces resulting HCF loss profile

HCF

VIAVI Solutions

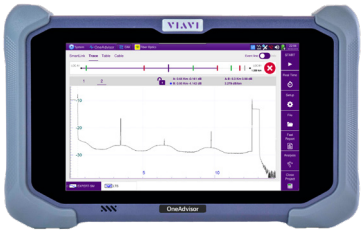
al deployments, Douglas Clague, solutions marketing
t time will protect rollouts from costly field troubleshooting

splices demand higher dynamic range at short pulse widths to resolve closely spaced events without dead zone merging.

Crest of a wave

Wavelength matters too. HCF is optimised for the S-band (from around 1450nm and above), so legacy 1310nm OTDRs won't be suitable. Testing at 1550nm is required, with 1625nm recommended for additional bend detection during field acceptance.

Given the transmission wavelength optimisation, attenuation profile (AP) testing identifies non-ideal transmission windows. This is important as HCF manufacturing processes, and resulting profiles, still vary between vendors.



Despite HCF's inherently low chromatic dispersion (CD) and polarisation mode dispersion (PMD), baseline dispersion characterisation remains valuable for validating installation quality and supporting future service upgrades. Be aware, OTDR derived dispersion measurements are constrained by low backscatter, so dedicated optical source and receiver based measurements are better suited to medium and long HCF spans.

Strategic direction

A credible HCF certification strategy combines bidirectional OTDR analysis for a true loss profile with careful measurement of fibre and splice losses, dispersion characterisation for a performance baseline and wideband spectral attenuation measurements to validate transmission windows.

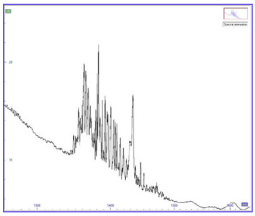


Figure 3: **Hollow core fibre attenuation profile measurements**

On the equipment side, OTDRs such as VIAVI's 8100 series D module deliver the performance HCF testing demands, with Report PRO software generating the HCF loss profile. Non-OTDR based optical broadband sources and receivers, like VIAVI's OBS-500 and Optical Dispersion Measurement module, allow AP measurement and dispersion testing over longer distances, all hosted on the OneAdvisor 800 Fiber platform. Furthermore, automated inspection microscopes, such as VIAVI's INX 700 with AI-enhanced image analysis, ensure fibre end-face cleanliness.

Next step

For more information on the unique challenges, test equipment requirements and best practice for HCF validation, **CLICK HERE** to download VIAVI's free application note.

www.viavisolutions.com

Onwards and upwards

In the first of a two part article on the past, prese and future of pre-terminated cabling, Jason Jam HellermannTyton looks at why it became essenti

As AI clusters have grown from hundreds to tens of thousands of graphics processing units (GPUs), the physical network infrastructure has become a significant engineering challenge in its own right. While much attention is given to switch application specific integrated circuits (ASICs), optics and transport protocols, the reality is that none of these technologies can function without a corresponding physical cabling infrastructure capable of supporting them at scale.

IN THE BEGINNING

In early data centres, structured cabling was relatively straightforward. Individual optical fibre pairs or small numbers of connections could be installed, terminated, tested and documented by engineers on-site. This approach worked well for traditional enterprise environments where network growth was gradual and port densities remained manageable.

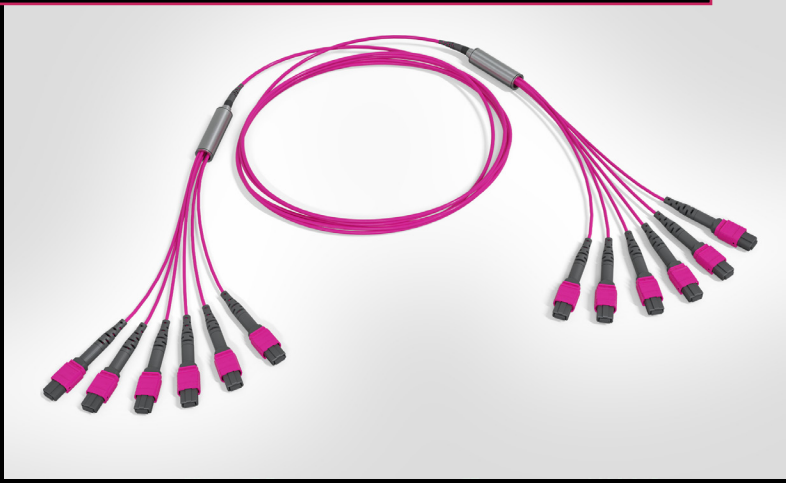


Modern AI factories operate at an entirely different scale though. Consider a cluster containing 2,000 GPU servers, each equipped with eight GPUs and their corresponding 400Gb/s network interface. Even before considering redundant paths or multi-plane architectures, the deployment requires approximately 16,000 individual network connections between servers and leaf switches. Once uplinks between leaf and spine switches are included, total fibre counts can easily exceed tens of thousands of links within a single fabric.

SUCCESS STORY

At the largest scale, AI deployments can involve hundreds of thousands of individual fibre connections, making the challenge far more complex than simply connecting devices. Success depends on achieving fast installation, robust quality assurance, consistent optical performance, effective cable management, long-term operational maintainability and, ultimately, reducing overall time to deployment. Traditional field termination methods struggle to meet these requirements economically or operationally.

The industry's answer has been the widespread adoption of factory terminated cabling systems. Rather than installing



individual fibre strands and terminating connectors on-site, cable assemblies are pre-manufactured, tested, labelled and certified in controlled factory environments before arriving at the data centre. Typical examples include MPO/MTP trunk cables, multi-fibre ribbon assemblies, structured spine to leaf cable harnesses and preconfigured rack-to-rack fibre bundles.

These assemblies arrive ready for installation with known optical characteristics and predictable performance, offering substantial benefits through factory controlled manufacturing. This approach enables lower insertion loss variation, improved connector consistency, reduced contamination risk, faster deployment, predictable quality control, simplified troubleshooting and greater subcontractor neutrality.

Most importantly, installation times can be reduced dramatically compared to traditional field terminated approaches. For a hyperscale AI cluster deployment, reducing installation effort by thousands of labour hours can significantly accelerate time to production.

SOLID FOUNDATIONS

An often overlooked benefit of pre-terminated infrastructure is that it naturally aligns with the industry's transition toward higher speed optics. As interface speeds increased, transceivers relied upon parallel optical lanes transmitting across multiple fibres simultaneously. An 800Gb/s optical module may contain eight optical lanes, while emerging 1.6Tb/s technologies may contain 16 or more lanes.

These architectures are inherently well suited to MPO based structured cabling systems because the required fibre groups already exist within a common pre-terminated assembly. In other words, the infrastructure designed to reduce cabling complexity also became the ideal physical foundation for next generation optical technologies.

> 'Perhaps the most important consequence of pre-terminated cabling is that it encouraged architects to think in terms of fibre groups and lane groups rather than individual links.'

CREATING THE CONDITIONS

Perhaps the most important consequence of pre-terminated cabling is that it encourages architects to think in terms of fibre groups and lane groups rather than individual links. Once multiple optical lanes are carried together within common cable assemblies, it becomes increasingly

practical to redistribute those lanes in different ways. This has led to the use of breakout or fanout configurations, where a single high speed optical interface can be divided into multiple independent lower speed connections.

This concept has been around for quite some time and began with 1x40G port = 4x10G ports. As expected, the technology didn't stand still and it is now viable to have a 1x800G port shared to 8x100G ports, 4x200G ports or 2x400G ports.

The physical infrastructure was already designed to transport groups of fibres, making these fanout models a natural extension of the cabling architecture.

Viewed through this lens, pre-terminated cabling was more than an operational improvement – it represented the first step toward the highly parallel networking designs now used in modern AI fabrics.

The progression from bundled fibre assemblies to optical lane fanout, to shuffle networks and, ultimately, to multipath reliable connection (MRC) can be viewed as a continuous evolution. Each stage builds upon the concept of treating connectivity not as a single dedicated path but as a collection of coordinated links that can be distributed, balanced and managed collectively.

GET TO THE POINT

Early high performance computing and AI systems relied on straightforward point-to-point cabling. Every server port was connected directly to a corresponding port on a network switch. In a typical deployment, a GPU server might expose eight 400Gb/s network interfaces. Each interface would be connected using a dedicated transceiver and fibre link directly into a leaf switch.

While simple to understand, this approach becomes increasingly problematic as clusters scale. A cluster containing 1,024 GPUs, eight network connections per server and 400Gb/s per connection already requires thousands of optical links and transceivers. At hyperscale AI deployments containing tens of thousands of GPUs, cabling density rapidly becomes a significant engineering challenge.



The result is greater installation complexity, higher optical costs, reduced airflow efficiency and increased risk of operational errors during deployment and maintenance. The industry therefore began looking for ways to

increase bandwidth without proportionally increasing cable counts.

KEEPING PACE

As Ethernet and InfiniBand technologies advanced, transceiver speeds continued to increase, progressing from 100Gb/s and 200Gb/s optics to 400Gb/s, 800Gb/s and emerging 1.6 Tb/s solutions. However, network switch and server interface technologies rarely evolved at the same pace, creating situations where a high speed optical interface needed to communicate with multiple lower speed ports. With breakout or fanout cabling an 800Gb/s transceiver could be logically divided into multiple independent links.

As discussed above, this concept has been around for quite some time and began with 1x40G port equalling 4x10G ports. This soon escalated to 1x100G port equalling 4x25G ports, before jumping again to 1x400G port equalling 4x100G ports. More recently a 1x800G port can become 8x100G ports, 4x200G ports or 2x400G ports.

BENEFIT CHECK

From the application's perspective, these links appear as independent network interfaces while physically sharing a common optical module. This innovation delivered several key benefits including reduced optical component counts, lower deployment costs, improved port utilisation and simplified migration between different technology generations. More importantly, it introduced the concept that a single high bandwidth endpoint could simultaneously participate in multiple independent network paths or planes.

Original lane swapping technologies used simple fanout cable assemblies to split the MPO interfaces of the faster ports into the duplex LC interfaces of the slower ports. Eventually, both splitting and combining functions needed to be combined and the shuffle cassette was conceived.

NEXT STEPS

This seemingly simple capability would become foundational for future AI network architectures and that's what my next article will look at. ■



JASON JAMES

Jason James is technical director at HellermannTyton. With over 35 years' experience and expertise in the LAN, data centre and fibre network infrastructure industry, James leads product development and new product introduction for the company. He also delivers technical training and thought leadership to professionals within the data centre and fibre broadband sectors.

The best of both worlds

Christophe Hinet of Molex Connected Enterprise Solutions highlights the benefits of adopting a hybrid approach to pre-terminated and field terminated solutions

Pre-terminated cabling is commonly presented with a simple proposition – reduce work and risk on-site by using pre-terminated lengths. Fewer field terminations and less time spent preparing connectors means a faster path from installation to operational link. Those advantages are real, but they are not the most important change.

ASSESSING THE BENEFITS

Pre-termination moves work upstream. Decisions that could once be resolved during installation have to be made earlier, often before an assembly is manufactured. Lengths, routes, connector types, port allocations, equipment positions, labels and delivery sequence all become part of the specification.

That can make an installation much more predictable but it can also make poor information more expensive. The useful question is therefore not whether pre-terminated cabling is better than field termination. It is which parts of a project are predictable enough to benefit from it and which still need the flexibility that field termination provides.

THE REAL STARTING POINT

Pre-termination is at its strongest where the physical installation is well understood, where cabinet locations are fixed, pathways are known, distances can be measured with confidence and large numbers of similar links need to be repeated. In those conditions, moving termination into a controlled manufacturing environment can remove a significant amount of repetitive work from the site programme.

Assemblies can arrive terminated, labelled and assigned to defined locations, leaving installers to focus on routing, inspection, cleaning if appropriate, connection and testing. There are other practical benefits too. In an occupied building, secure area or tightly managed construction environment, reducing cutting, preparation, waste and specialist termination work on-site can simplify access and sequencing. The benefit is not only speed but also greater control over where the work is performed.

HORSES FOR COURSES

However, not every project offers that level of certainty and refurbishment is an obvious example. Existing containment may differ from drawings, access may be more restricted than expected and equipment positions can move once walls, ceilings or cabinets are opened up. Even on a new build, a late change to a cabinet position or pathway can alter the required length of many links.

Field termination retains value in precisely those circumstances. It gives the installer room to respond to what is present rather than what was expected to be present. For many projects, the right answer is therefore hybrid – pre-terminate the high volume, repeatable links and



retain field termination where uncertainty remains.

FIT FOR PURPOSE

Factory fitted connectors do not remove the physical constraints of a copper installation. A pre-terminated copper assembly still has to pass through containment, negotiate bends, enter cabinets and leave a manageable amount of cable at the destination. The route needs to be considered as an installed path rather than a straight line distance on a drawing.

Length management is especially important. A cable that is too short is clearly unusable, but adding generous spare length to every assembly is not a good substitute for accurate planning, and excess cable still needs to be stored and managed. In dense pathways or cabinets, unnecessary service loops can consume space, obstruct access and make future moves and changes harder.

Where balanced cabling will also support remote powering, installation conditions matter for another reason. Cable grouping, pathway conditions and the surrounding environment can influence thermal behaviour. Pre-termination changes where the connector is fitted – it does not remove the need to assess how the installed cabling will operate. The objective should not be to order an assembly that merely reaches from A to B. It should arrive at the correct length, follow the intended route and remain manageable once the installation is complete.

LIGHT WORK

With pre-terminated fibre optic cabling some of the priorities change, with connector condition an obvious example. A factory terminated connector may leave manufacturing in excellent condition but it still has to be packaged, transported, unpacked, routed and connected.

‘Pre-terminated cabling is usually better for predictable, repeatable installations, while field termination is better where site conditions or lengths remain uncertain.’





Protective caps and packaging reduce exposure – they do not make inspection and appropriate cleaning unnecessary at the point of installation.

Polarity and port mapping become increasingly important as density increases. MPO/MTP based connectivity can combine multiple fibres in a single interface, so an error in mapping can affect several channels at once. A cable can be physically intact, correctly installed and still not provide the intended end-to-end connectivity if polarity or port assignments are wrong. That makes documentation operational, not administrative.

Factory testing is equally valuable but it is only one stage of verification. An installed assembly has been handled, routed and mated in conditions that cannot be assessed before it reaches site. Connector inspection and the project’s required performance testing still belong in the installation and acceptance process.

SAME PLACE SAME TIME

The cable and the information must arrive together and this is where pre-terminated projects often succeed or fail. A cable schedule should describe the real installation, not just provide enough information to place an order. Measurements need to reflect the pathway, changes in elevation, cabinet entry points and the service length genuinely required. Identifiers should remain consistent between drawings, labels, packaging and test records.

The same discipline applies to logistics. If hundreds of individually specified assemblies arrive at once with no clear relationship to the installation sequence, termination work may have been removed, only to create a sorting problem instead. Packaging should help the site team answer three questions immediately – where does this go, which end is which and when is it needed? A good pre-terminated system reduces decisions at the point of installation. It should not create a second design exercise on the floor.

Change control becomes part of cabling



design too. Once a fixed length assembly has entered manufacture, moving a cabinet or changing a route can have a direct physical consequence. Projects therefore need a clear point at which information is considered reliable enough to release, together with a process for communicating changes that occur afterwards.

That is the trade-off at the centre of pre-termination. The installation gains certainty because more work has already been completed, but the project must provide greater certainty before that work begins.

RIGHT FIRST TIME

Pre-terminated cabling should not be an objective in itself. Most links can be supplied with connectors already installed but the more useful question is whether doing so removes enough complexity from sites to justify fixing more of the design earlier.

For repeatable, well defined links, it often does. Work can be completed under controlled conditions, site activity becomes more predictable and installation teams receive assemblies prepared for a known purpose. Where routes, equipment positions or final requirements are still fluid, field termination may remain the more practical option. That is not a compromise or a failure to modernise the installation, it is the appropriate response to uncertainty.

VALUE ADD

The most effective projects recognise both conditions and design accordingly. Seen in that way, pre-termination is less a different type of cable than a different way to organise the work. Its greatest value comes not from trying to eliminate field termination everywhere but from moving predictable work off-site while preserving flexibility where the project genuinely needs it. Put simply, pre-terminated cabling is usually better for predictable, repeatable installations, while field termination is better where site conditions or lengths remain uncertain. ■



CHRISTOPHE HINET

Christophe Hinet is director of technical support at Molex Connected Enterprise Solutions, leading global technical support and cross functional collaboration. With 33 years in structured cabling, he bridges technical, commercial and strategic insight, advocates proactive pre-sales, drives customer centric solutions and champions innovation through practical digital tools worldwide for enterprise building and data centre customers. A recognised contributor to the industry, he holds the BICSI RCDD credential.

Molex Connected Enterprise Solutions

Molex Connected Enterprise Solutions provides end-to-end copper and fibre optic structured cabling for enterprise buildings and data centres. Backed by the global engineering and manufacturing capabilities of Molex, our portfolio includes cabling, connectivity and cable management, with solutions for commercial buildings, healthcare, education and other enterprise environments.



As a partner for consultants, integrators and installers, we support projects from initial planning through to installation and testing. With custom configurable pre-terminated assemblies, modular fibre optic systems and zone cabling, we offer options for new builds, expansions and upgrades to existing installations.

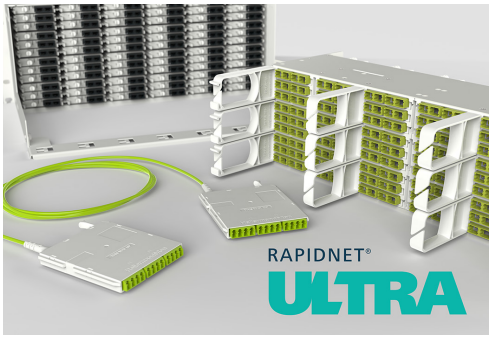
Customers have access to technical specialists, design resources, product documentation and planning tools to help design, build and maintain your infrastructure.

Through our global network of sales and technical teams, distributors and certified installers, we connect customers with solutions and support to help plan and deliver enterprise infrastructure that meets your application and operational needs.

CLICK HERE to find out more.
www.molexces.com

HellermannTyton

HellermannTyton is a pioneer in pre-terminated infrastructure with its RapidNet solution. RapidNet is the pre-terminated, pre-tested cassette based plug and play system developed by HellermannTyton that is designed to deliver high performance and improve the ease and speed of infrastructure deployments.



HellermannTyton's RapidNet ULTRA takes the existing data centre optical fibre solution beyond today's requirement, offering an even greater fibre density while accommodating very small form factor (VSFF) connectivity. RapidNet ULTRA supports tomorrow's requirements for high bandwidth, advanced network architectures and Ultra Ethernet.

With a range of cassette formats and fibre assemblies, RapidNet ULTRA delivers high performance and fibre capacity to meet the demands of the modern day data centre.

To find out more **CLICK HERE.**
www.htdata.co.uk

Mayflex

Save time and reduce costs with Mayflex Config Services

If you're installing Excel Networking Solutions copper or optical fibre cabling and haven't explored our Config Services, you could be missing a simple way to save both time and money.



Our Config Services facility provides pre-terminated copper, fibre and MPO solutions manufactured to your exact project requirements. Every assembly is produced in a dedicated clean, controlled environment using the latest termination and testing equipment, ensuring consistent quality and performance.

Each connection is individually inspected and tested before leaving our facility, reducing the risk of on-site issues and eliminating the time associated with field termination and testing. From pre-terminated fibre trunks and MPO solutions to customised copper assemblies, we can significantly reduce installation time, accelerate project delivery and free-up valuable engineering resource.

Less time on-site. Less waste. Faster deployment.

CLICK HERE to find out more about Mayflex Config Services.

www.mayflex.com

MISSED AN ISSUE?

CLICK ON THE COVER TO READ MORE



FOR A FREE SUBSCRIPTION TO **Inside_Networks** CLICK HERE

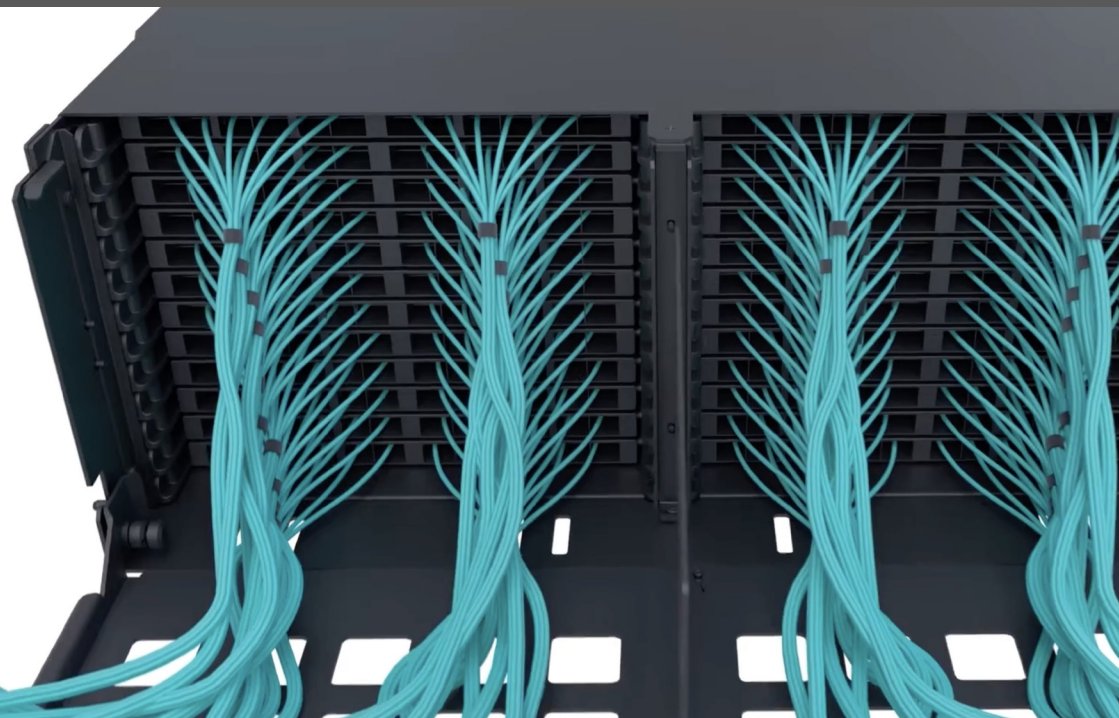
A question of time

Michael Akinla of Panduit explains why pre-terminated cabling is gaining ground and the key considerations when specifying and installing optical fibre and copper variants

The most expensive part of a cabling installation is not always the cable. In a live data centre, a hospital or a busy commercial building, it may be the time spent working around other trades, securing access and waiting for a new connection to be ready. Pre-terminated copper and optical fibre systems shift much of the preparation into a controlled factory environment. Cables arrive cut to a specified length, terminated and tested, ready to be routed and connected on-site. That change can make a substantial difference, provided the design work happens early enough.

TIME BACK ON-SITE

The principle applies to both media, although the assemblies look different. Copper will arrive as a cable loom of six, typically with jacks, modular plugs or cassettes at one or both ends. Fibre may use trunks, harnesses, interconnects and cassettes fitted with LC or multi-fibre connectors. An installer still needs to plan pathways, protect connectors during the pull, dress the cable, label the ports and verify the installed link. The attraction is that repetitive termination work no longer competes for time and space on the project floor.



Factory termination is particularly valuable where dozens or hundreds of similar links must be installed to a deadline. Instead of making each connection at a rack or outlet, the team can route a completed assembly and mate it with the appropriate panel or cassette. This reduces the specialist work required in a restricted space and makes labour demand more predictable. A phased installation also becomes easier to coordinate – cable pathways and cabinets can be prepared ahead of an equipment delivery, with final connections made during a short access window.

CHECK MATE

Testing is another advantage, but a distinction matters. Factory test records provide evidence of an assembly's performance before shipment. They do not prove that a cable has survived transport, pulling, bending and installation without damage. Site inspection and the testing required by the project specification remain essential. The strongest approach combines documented factory quality with a practical acceptance process after installation, rather than treating a pre-terminated assembly as a link that needs no checks.



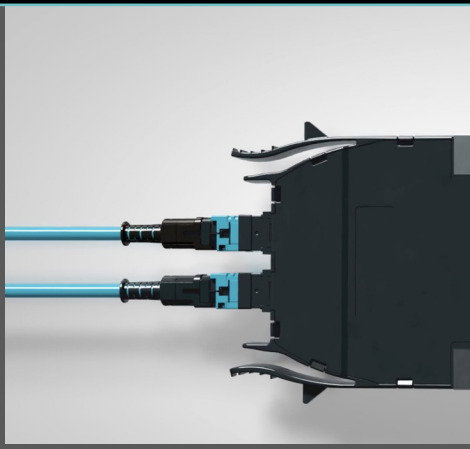
Fibre especially benefits from moving connector preparation away from site. Cleaning, polishing and testing optical connections demand care, and crowded or dusty environments are rarely ideal. Factory made multifibre trunks can also provide substantial capacity in a compact pathway. However, an optical design still needs a loss budget covering every mated connection. Connector type, polarity, fibre count and cassette arrangement must match the active equipment and planned upgrades. A trunk that fits mechanically is not necessarily the right optical channel.

Copper has its own design questions. Category rating, shielding, connector configuration and the intended application need to be specified together. Power over Ethernet (PoE) loads may influence cable selection, bundle size and heat management. The route must accommodate the connector ends and any pulling protection without exceeding the permissible bend radius or pull tension. Ordering a bundle of links at the wrong length can turn labour saving into congestion in the cabinet. Good measurements, an agreed allowance for routing and clear identification at both ends are therefore fundamental.

> 'There will always be jobs where field termination offers more flexibility. Yet when routes are known, links repeat and access is precious, pre-terminated copper and fibre turn cabling into a more planned construction activity.'

EARNING ITS PLACE

For hyperscale, colocation and enterprise data centres, repeatability is the compelling factor. Operators need to bring capacity online in stages while preserving a clear



record of which ports serve which systems.

Pre-terminated trunks and modular distribution hardware support a standard rack design that can be repeated across rows or sites. Compact fibre assemblies help where port density is high, while copper trunks can simplify the many shorter connections between distribution points and equipment. Faster deployment matters, but orderly moves, adds and changes (MACs) are just as important once a facility is operating.

LIVE AND DIRECT

Multi-tenant operators have a related concern – every customer fitout takes place in a live environment. A planned route and a tested assembly can reduce the work done near existing services and help limit the time contractors spend in controlled areas. The benefit is operational as well as commercial. If a customer needs a connection on a specific date, uncertainty about on-site termination and rework can affect handover.

Hospitals, research facilities and other critical environments often favour the same predictability for different reasons. Access to technical rooms may be tightly managed and MACs can be difficult to schedule. Standardised assemblies allow the installer to prepare carefully and carry out a shorter, more controlled change.

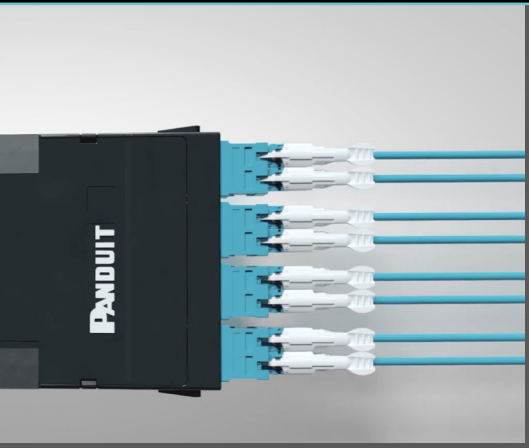
A Panduit project at a Luxembourg medical research institute, for example, used pre-assembled copper and fibre cabling as part of a redundant data centre build designed for future reconfiguration. The broader lesson is that installation speed has value when it reduces disruption to an organisation whose work cannot pause for a cabling project.

REPEAT ORDER

Commercial campuses, logistics sites and industrial facilities can benefit where their network builds are repetitive. Rows of wireless access points, cameras, control cabinets or equipment zones may follow a common layout.

Pre-terminated copper can speed a planned rollout, particularly when the endpoint positions are fixed and the routes have been surveyed. Fibre trunks can serve longer backbone routes and distribution points. For distributed sites, repeatable assemblies also help a central team maintain a consistent design across multiple contractors and locations. Where layouts change frequently or routes are uncertain, some field termination may





still be necessary.

Broadband providers and contractors face another version of the time problem. Pre-connected fibre drops and service terminals can reduce field splicing and the number of specialist visits needed to connect premises. Their outdoor design must suit the actual route and exposure – a data centre trunk is not an interchangeable substitute. For the service provider, the goal is to connect more customers with available crews while leaving a network that can be maintained when the original installation team has moved on.

SPECIFY THE WHOLE

The decision to use pre-terminated cabling should factor in installed cost and operational risk, not the purchase price of cable alone. Assemblies require accurate schedules, agreed connector and polarity conventions, lead times and a method for handling late design changes.

Projects should confirm applicable European cable fire performance requirements for the installation and the documentation supplied for the selected cable. Packaging, spare lengths and the potential to reuse modular components should also be included in the material plan. Fewer offcuts and less site work are plausible benefits but they depend on ordering the right lengths and protecting assemblies, so they do not have to be replaced.

COMMAND AND CONTROL

There will always be jobs where field termination offers more flexibility. Yet when routes are known, links repeat and access is precious, pre-terminated copper and fibre turn cabling into a more planned construction activity. The real gain is not simply connecting a panel faster, it is giving designers, installers and operators more control over quality, schedule and the next network change. ■



MICHAEL AKINLA

Michael Akinla is business manager central Europe north at Panduit. He brings over 20 years' experience in the deployment of Panduit's most complex solutions and has extensive experience in working with several large global accounts to bring about significant improvements in terms of higher bandwidth deployments, reduced Power Usage Effectiveness (PUE) ratings and lower total cost of ownership.

Do the shuffle

In this month's Knowledge Bank, where tech is demystified, Carrie Goetz examines optical fibre configurations for AI

What used to be a popular dance is now an AI strategy. The fibre configuration for a graphics processing unit (GPU) cluster is vastly different than a fibre configuration for a 100Gb/s network. When someone submits a prompt to AI the crunching begins, but when you consider the vast amounts of data, 100Gb/s is not enough. Moving from that network to a GPU 800Gb/s or the anticipated 1.6Tb/s networks isn't as simple as changing out a server or a patch cord though.

At the rack, pathways and bend radii change with the addition of liquid cooling. Cable trays need more clearance for cooling loop manifolds and maintenance. Pathways will likely need to be replaced or supplemented with additional tray and we are no longer anticipating all ports connecting in the rear of a server. In these cases, very small form factor (VSFF) connectors, or the network cabling, is no longer in the rear of the cabinet. The trend is for these connections to sit in front of the server, unobstructed and, unfortunately, away from the rear aligned basket trays. Power may also move, while piping clearances and pathways need to be defined too.

Beyond the physical rack changes, network connections from GPU use 800Gb/s octal small form factor pluggable (OSFP) transceivers. The network transmits four lanes of 200Gb/s or eight lanes of 100Gb/s over ultra-low loss fibre OS2 trunks with MPO-16, dual MPO-12 or MMC16/24 very small form factor (VSFF) connectors. The layout of the fibre plant is set in stone prior to AI installations and will remain unchained. Let's explain why.

Beyond the location of the connections,



how the light travels is a bit different than we are used to seeing for lower speed applications. AI uses fibre multi-lane processing. While we have used parallel processing for some time now for Ethernet and Fibre Channel, for AI applications the configuration, yet again, changes. In parallel processing data moves along multiple paths (strands) from point A to point B



simultaneously. For AI, however, multiple conversations and communications fire everywhere, always, all at once.

GPU shelves connect to multiple networks, while ports connect to other GPUs in a mesh fashion. But how do we get all the parts communicating in concert? We let them have a shuffle party.

Traditionally, a server connects to a single switch and the server communicates with the other devices on the network through the switch port. The switch decides how to move the packets. To do so, the switch needs to look at the packet being transmitted, identify the destination address, figure out which port (local or uplink) it is attached to and route it accordingly. At Layer 2 the switch looks at the MAC address and at Layer 3 the switch looks at best routing information.



AI processes significant amounts of information at once and communicates with multiple networks at the same time. Communicating over multiple fibres simultaneously via a cluster increases functionality, agility and speed. One way to accomplish this is a shuffle. A shuffle is passive and is the equivalent of opening the floodgates with nobody minding your business. It examines nothing, as it is passive, but rather serves as a mapping mechanism for fibre allowing communications across and to multiple switch planes (ports).

The magic happens in the shuffle cassette (or in some cases trunks). The shuffle cassette fans the incoming fibres out to multiple switch ports/planes. An AI switch has a traditional row of dense OSFP (or QSFP-DD) cages for small metal transceivers with a single MPO-16 connector, two MPO-12 connectors or the new higher density miniature duplex connector (MDC) or Senko Nano (SN) connectors which are half the width of a traditional LC.

Typical shuffles are 2x2, 4x4 or 8x8 formats. For instance, a 4x4 takes four MPO assemblies from four disparate GPUs into the shuffle cassette and then internally splits the fibres for connection to four switch ports/planes. Any one server/GPU port talks to four switch ports at once. The cassette's outgoing shuffle ports now contain light mapped from each of the four servers. The real beauty here is the speed and spread in which the data moves. It would be the equivalent of being able to have four people talk to you at the same time with full grasp of all they are saying, in real time.

Depending on the actual tasks at hand, the lanes can be dedicated to networking or storage, or there can be any variety of custom mappings for various AI services. One thing is certain – this type of fibre needs a map that is fixed and part of the overall official documentation. Testers will require more than a power meter and light source and test lane mapping. Due to the sheer volume of fibres and the intricacy of the designs, some AI expertise will likely be needed during the design phase. ■

Quick clicks

Your one click guide to the very best industry events, webinars, electronic literature, white papers, blogs and videos

Building The Right Foundation For Liquid Cooling Design is a white paper co-authored by **LiquidStack** and **Leading Edge Design Group (LEDG),** based on their comparison of notes from real world liquid cooling deployments.
CLICK HERE to download a copy.

SolarEdge and **NVIDIA** have published a joint white paper titled 800VDC Protection And Grounding For AI Data Centers.
CLICK HERE to obtain a copy.

Developing And Protecting Critical National Infrastructure For The UK's Digital Future is a white paper from **Altnets** that examines how the UK can safeguard its critical national infrastructure through resilience, physical security and data sovereignty.
CLICK HERE to download a copy.



FOR A FREE SUBSCRIPTION TO **Inside_Networks** CLICK HERE

Supporting Smarter Infrastructure Design Decisions In Smart Buildings is a blog by Curtis Miller and Randy Nottingham of **Siemon.**
CLICK HERE to read it.

A DCIM Playbook To Close The AI Readiness Gap is an ebook by **Sunbird.**
CLICK HERE to request a copy.



Turn Link-Loss Budgets Into Operating Discipline is a practical guide from **Molex Connected Enterprise Solutions.**
CLICK HERE to get a copy.

Designing Cabinet Infrastructure For The Next Wave of AI-Driven Hyperscale Density is a blog by Maren Price of **Chatsworth Products (CPI).**
CLICK HERE to read it.

The building blocks of success

Ben Pritchard of AVK explains why the future of AI power is modular

The industrialisation of AI has begun and this, in turn, is accelerating worldwide data centre infrastructure rollout. Together the combination of these forces is shifting every data centre development dynamic and this is particularly true for power provision.

CHANGING TIMES

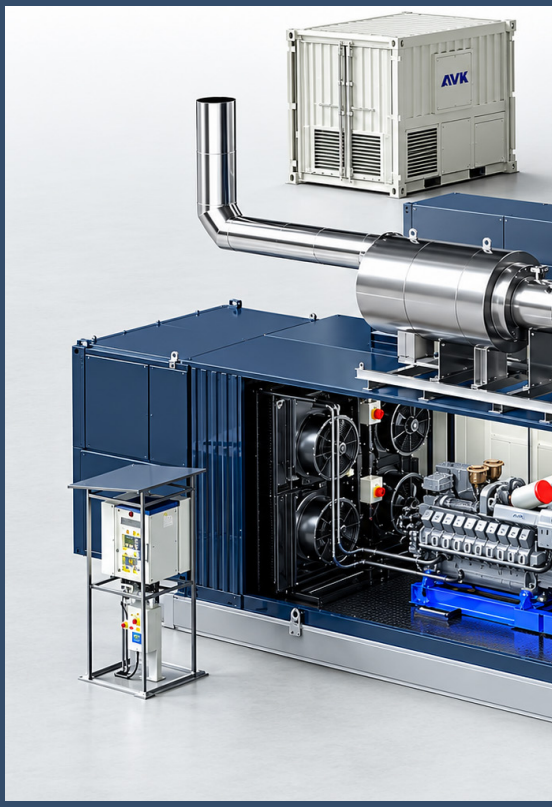
In the pre-AI era, the task list for data centre developers usually began with finding and acquiring land with available power or permission to build a substation. Today, project starting points are both simpler and more complex.

Where the powering of many existing and emerging types of AI is the now primary objective for developers, the challenge is finding enough power to match the AI workload. We're entering a new generation of data centre infrastructure types with power capacity needs ranging from GW scale AI factories to traditional shed form factor facilities running tens and hundreds of MWs, and bespoke inference workload AI edge data centres dotted around population centres, financial districts or business and research clusters.

POWER AT SPEED

This has radically altered the landscape to one where modularity is integral to the data centre build environment. Innovation in power solutions to meet the speed and scale demands of developers is the new game in town. Because of this, factory engineered, fully integrated modular power units are a vital component of this new value chain.

Whatever the data centre development use case, in new builds, site upgrades and repurposing of pre-existing facilities, AI factories are being built through adoption of pre-fabricated construction modules comprised of containerised power, modular buildings and factory built AI halls. The



'As developers turn to factory built power infrastructure, modular power units deliver exactly what AI factory development demands – speed, cost efficiency and reliability at scale.'

market is expanding into diversified data centre types across commercial data centres, hyperscalers and the new breed of edge and neocloud operators building to accommodate graphics processing units as a service (GPUaaS) for AI at scale. However, all data centre builders are faced with the same macro trends:

- Major issues with power access, security and sustainability.
- Lengthening equipment supply chains for generators, switchgear, uninterruptible power supplies (UPS) and storage.
- Lengthening grid connection queues.
- Strained grids – requiring decade long upgrade projects.
- End customers who can't wait.



MEETING THE NEED

There are several ways that modular power units deliver:

- **Power delivery built around infrastructure repeatability**
  As developers turn to factory built power infrastructure, modular power units deliver exactly what AI factory development demands – speed, cost efficiency and reliability at scale.

- **Capital efficiency**
  Modular power unit scaling is financially efficient for power as you grow provision. Developer and operator provision of customer matching optionality to deploy a 2MW, 5MW or 10MW unit as demand lands is key. This also radically improves capital efficiency for every type of data centre owner.

- **Speed to power deployment**
  As AI scales the power challenges presented to operators are not confined to densities. For corporate users of AI infrastructure and GPUaaS providers there is a new pressure in the form of the AI execution gap.
  Whereas previously, corporates wouldwork to 12-18 month timeframes for major software changes or digital transformation projects, in some

cases today's AI is expected to be operational in as little as three months. In terms of IT development this is known as the 90 day production sprint. This is where modular power solution buildouts can respond to rapid ready for service timeline pressures of AI workloads.

- Avoidance of supply chain bottle necks

  Lead times are increasing for data centre equipment such as diesel generators, power distribution systems, medium voltage transformers and advanced cooling systems, with delivery sometimes taking several years.

TALKING POINT

The issue of AI, data and digital sovereignty is a top table discussion for every government and business across Europe. The moves to sovereign AI and power procurement have already begun and an interesting consideration here is that of supply chain sovereignty.

Data centre operators need ultimate control over the components and infrastructure utilised within the modular power units that comprise their critical electrical infrastructure. In short, sovereignty isn't only about where data sits, it's about who controls the physical infrastructure keeping it running.

The data centre industry is facing a power timeline disconnect between securing power infrastructure and the market requirement for power where it is needed at speed. This is because the supply chain timelines for equipment are lengthening at exactly the moment when the need for speed to power is top of the list of operator customer concerns.



These concerns cut across the commercial data centre developers whose customers are AI companies, enterprises building private and hybrid AI clouds to the global AI large language model (LLM) training firms, neoclouds and the public cloud firms building AI factories at GW scale. All are seeking ready for service certainty for customer commissioning and energisation timeline certainty. But addressing the speed to power timeline certainty begins long before customer rollouts.

FORM AN ORDERLY QUEUE

Across Europe, getting on to grid connection queues is a time consuming and increasingly uncertain process that can materially impact planning and permitting timeframes. Markets are becoming



increasingly complex, making it more difficult to ensure energy solutions meet bespoke needs for workload performance, resilience, cost and flexibility combined with broader environmental responsibility.

The same applies to power distribution for increasingly complex AI technologies, where traditional power distribution approaches struggle to match high density compute specialised and failover strategies. Cutting edge modular power units are built out using standardised components without losing flexibility and customisation options. This enables design responsiveness for scale out and equipment refreshes that match power density and cooling density AI accelerator roadmaps.

TIME IS OF THE ESSENCE

Whatever its headline MW figure, no data centre has ever reached maximum capacity on day one. The old model of build 100MW and wait for the available power to reach utilisation is no longer fit for purpose for many AI operator use cases. Although not obvious to the untrained eye to date, almost every data centre ever constructed has an individual uniqueness to its power chain topology. Modular data centre architecture will not replace how every data centre is powered but for a growing number of use cases in existing and new environments, modular power units are being recognised as the perfect fit. ■



BEN PRITCHARD

Ben Pritchard is CEO of AVK. With over a decade of experience managing and delivering complete end-to-end critical power projects, numerous hyperscale data centre owners and operators trust Pritchard's experience and extensive knowledge of energy technology to successfully deliver multimillion pound contracts.

Delivering the goods

Pierre Ecuyer of Legrand explains why AI is redefining uninterruptible power supply (UPS) resilience

AI is changing the data centre industry in ways that go far beyond increases in rack density. New AI and high performance computing environments are placing greater demands on every layer of the electrical chain, from grid connections, transformers and switchgear to batteries and power distribution systems – including UPS systems.

TWICE AS NICE

While the UPS has traditionally been viewed primarily as protection against power outages, its role in maintaining electrical stability is becoming increasingly important. Graphics processing unit (GPU) clusters can create highly variable and unpredictable changes in power demand.

When large numbers of processors are synchronised, changes in demand can occur at the same time, creating steep ramp rates, large step loads and greater power quality challenges. Next generation UPS systems therefore have a dual role to play – protecting sensitive IT equipment from disturbances in the incoming power supply while helping to manage dynamic AI loads and limit their impact on the wider electrical infrastructure.

SHIFT WORK

For data centre operators, this changes what resilience looks like. AI workloads can shift rapidly between periods of low utilisation and intensive processing, creating a much more variable electrical profile than traditional IT environments.

Critical power systems need to maintain stability as loads fluctuate, operate efficiently across a wide range of load conditions and scale as capacity requirements grow. Resilience is no longer just about what happens when power fails, it is also about how the power architecture performs when power is available but demand is changing rapidly.

A NEW PERSPECTIVE

This has implications for how operators, consultants and specifiers assess UPS systems. Autonomy time, nameplate capacity, redundancy and headline efficiency remain important but they no longer tell the whole story. Low load efficiency, dynamic load response, modular expansion, service continuity and losses across the operating curve all deserve greater attention.

In double conversion mode, the UPS is already doing more than waiting for an outage. It provides a power conditioning layer between the incoming supply and sensitive IT equipment, helping to isolate the load from disturbances and maintain stable power quality. As AI makes the load itself more dynamic, this aspect of UPS performance becomes much more significant.

THE REAL OPERATING CURVE

AI also changes the way UPS efficiency should be considered. Maximum efficiency

figures at or near full load provide a useful benchmark but they may not reflect actual operating conditions. AI infrastructure will not necessarily run at full utilisation all the time, making efficiency at partial load an important factor in energy consumption and operating costs.

> 'Advances in SiC technology, modular architectures and intelligent monitoring are enabling UPS systems to combine high efficiency and scalability with a more active role in managing critical power.'

Next generation UPS technology is responding to this challenge. Advances in power electronics, including silicon carbide (SiC) technology, can reduce electrical losses and support high efficiency across a broader range of operating conditions. Three phase modular UPS systems are now available that can achieve 98.6 per cent efficiency in voltage and frequency independent (VFI) double conversion mode at 25 per cent load and 98.4 per cent at 50 per cent load. In an AI environment, maintaining high efficiency across the operating range can be just as important as the peak figure.

There is also a thermal consequence. Energy lost during power conversion is released as heat and ultimately must be removed by the cooling system. With power and cooling capacity already at a premium in high density facilities, reducing those losses can deliver benefits on both sides of the equation.



BUILDING FLEXIBILITY

Predicting the power requirements of an AI deployment several years ahead is difficult. Operators need sufficient capacity to meet today's requirements but building additional capacity around uncertain future demand can leave expensive infrastructure underutilised.



Next generation modular UPS architectures offer a more flexible approach, allowing operators to add capacity as demand grows. Modularity can also make systems easier to maintain. Features such as hot swappable power modules and front access serviceability can enable maintenance, repair and expansion with less disruption to protected loads.

In a mission critical environment this flexibility is an important part of resilience. It is not simply a question of how much redundant capacity has been installed but how easily infrastructure can be maintained, expanded and adapted while remaining operational. For colocation providers this can help align investment more closely with customer demand. For enterprise and hyperscale operators it provides greater flexibility to respond as their AI requirements develop, rather than committing today's infrastructure to assumptions about future capacity.

A MORE ACTIVE ROLE

Another change is taking place beyond the AI load itself. Grid constraints, renewable integration and energy volatility are placing greater emphasis on how large data centres interact with the electrical network.

The primary purpose of the UPS remains protecting critical loads and ensuring

service continuity but advances in UPS technology are enabling it to contribute more to the wider power strategy. Advanced monitoring, intelligent battery management, smart grid readiness and capabilities such as peak shaving can give operators greater visibility and control over how power is used. Battery control also becomes more important as loads become more dynamic, balancing stored energy use to support electrical stability against unnecessary charge and discharge cycles that can contribute to battery wear.

That could become increasingly valuable as AI drives demand upwards while grid capacity remains constrained in many markets. The ability to manage power more intelligently is as important as securing additional capacity.

ONE STEP BEYOND

The rise of AI is forcing data centre decision makers to look beyond simple capacity planning. More megawatts will be required but designing for AI is not simply about adding capacity. Critical power infrastructure increasingly needs to be designed for load dynamics as well as the number of kilowatts it must support. Operators need critical power architectures that can respond to dynamic loads, remain efficient at partial load, scale as demand evolves and be maintained without unnecessary disruption.

Next generation UPS technology is evolving around these requirements. Advances in SiC technology, modular architectures and intelligent monitoring are enabling UPS systems to combine high efficiency and scalability with a more active role in managing critical power.

CHANGING TIMES

Ultimately, AI is not changing the UPS's fundamental purpose – protecting critical infrastructure and maintaining continuity remain its primary responsibilities. What AI is changing is the environment in which the UPS must fulfil those responsibilities. As AI adoption accelerates, power resilience will increasingly depend not only on what happens when power fails, but on how efficiently, reliably and intelligently the power architecture performs every moment the data centre is operating. ■



PIERRE ECUYER

Pierre Ecuyer is director of critical power and services at Legrand, where he leads strategic initiatives across UPS, transformers, high power busbars and service innovation. With over 25 years of experience in the data centre infrastructure industry, he brings deep expertise in scaling service operations and aligning technology with business outcomes

The Energy Group

Powering resilience in critical infrastructure

As businesses become increasingly dependent on uninterrupted power, resilience and energy efficiency have never been more important.

The Energy Group delivers comprehensive critical power and energy management solutions, helping organisations protect essential infrastructure, reduce operational risk and improve energy performance. Specialising in UPS, central battery systems, standby generators and preventative maintenance, The Energy Group supports businesses across data centres, telecommunications, commercial facilities and the public sector.

From initial power assessments and system design through to installation, commissioning and ongoing maintenance, our team provides end-to-end support tailored to each customer's requirements. Our wider portfolio includes solar PV, battery energy storage and intelligent energy management, enabling organisations to combine traditional power protection with modern energy technologies.



Whether upgrading existing infrastructure or planning future capacity, The Energy Group helps businesses build more resilient, efficient and reliable power systems.

CLICK HERE to find out more and to send us an email **CLICK HERE.**
www.theenergygroup.uk

Elevate

Elevate's new Intelligent Power Distribution Units (iPDUs) provide data centre operators with the visibility, control and flexibility required to manage modern IT environments more efficiently. Designed for enterprise, colocation and cloud applications, the range combines advanced monitoring, remote management and intelligent power distribution within a scalable platform.

A standout feature is the patented universal 4-way outlet design, supporting C13, C15, C19 and C21 connectors from a single socket. This simplifies deployment, reduces inventory complexity and provides greater flexibility when supporting mixed equipment environments.



The range delivers detailed power monitoring at both iPDU and outlet level, enabling real time visibility of energy consumption, capacity and utilisation. Hot swappable controllers help maintain uptime during maintenance, while remote switching and reboot capabilities support efficient management of connected equipment.

With support for centralised management of up to 40 iPDUs through a single software platform, daisy-chaining capabilities and stocked models available for rapid deployment, Elevate iPDUs offer a smarter approach to power management in today's data centre environments.

To learn more about the new Elevate iPDU range **CLICK HERE** or contact the Elevate team.
elevate.excel-networking.com

Mayflex

Delivering power protection for every requirement



Power interruptions can have a significant impact on networks, security systems and critical business operations. That's why choosing the right UPS solution is essential.

At Mayflex, we're Driven to Deliver More, offering a comprehensive range of power management solutions from both APC by Schneider Electric and Uniti Power. Together, these trusted brands enable us to meet a wide range of application, performance and budget requirements.

APC delivers industry leading power protection and advanced battery management, while Uniti provides reliable, cost effective UPS solutions with intelligent monitoring and flexible deployment options.

Our support goes beyond supplying products. The Mayflex team has the knowledge and expertise to help customers select the right solution for their environment, ensuring the best balance of protection, performance and value.

Whatever your power management needs, we're here to help keep your critical systems running.

To find out more **CLICK HERE.**

www.mayflex.com

Networks Centre and Legrand

Powering the neocloud opportunity together



Networks Centre and Legrand are working collaboratively to capture the growing neocloud market, delivering reliable, scalable UPS power protection solutions to meet the rapidly increasing demands of AI and high performance computing infrastructure.

With power resilience at the heart of every data centre, our joint focus is to ensure customers can access the right UPS solutions quickly, backed by the technical expertise, product availability and responsive service needed to keep projects moving.

By combining Legrand's Power Control UPS expertise with Networks Centre's agility, distribution capabilities and customer relationships, we aim to simplify procurement, accelerate project delivery and help neocloud providers deploy resilient power infrastructure at pace.

Together, Networks Centre and Legrand will bring dependable UPS solutions to market faster – giving customers the confidence, availability and support they need to scale.

CLICK HERE for more information.

www.networkscentre.com

Striking the right balance

Colin Dean of Socomec examines the role of uninterruptible power supplies (UPS) in the AI era and how data centre operators can combine resilience with efficiency

The boom of AI computing is reshaping the power requirements of digital infrastructure. Rack densities that once sat at 5-10kW are now exceeding 100kW as standard in high performance deployments, with some reaching even higher levels. What's more, AI presents more unpredictable power loads than traditional compute, driving surges in demand that test the limits of legacy infrastructure.

PRESSURE TEST

With the number of UK businesses deploying AI tripling over the past three years, these challenges are only growing in severity. Data centre operators are under immense pressure to ensure uninterrupted power for AI applications, while simultaneously facing growing scrutiny over the environmental footprint of their facilities.

This extends to stringent regulatory requirements including the European Union's Energy Efficiency Directive (EED), which mandates that facilities over 500kW submit their Power Usage Effectiveness (PUE) metrics on an annual basis and demonstrate year-on-year efficiency improvements. Operators in the UK are increasingly adopting EED style reporting as best practice, driven by the need for consistent design, comparable performance metrics and competitiveness across international markets.

One of the key areas in which data centre operators can make strides to improve the resilience and efficiency of their facilities is UPS systems. Originally designed to keep power flowing even in the event of a fault, these systems play in an increasingly important role in sustainably accommodating the demands of AI and, in turn, helping operators achieve the kind of efficiency gains that contribute to the overall competitiveness of their facilities.

COST OF REDUNDANCY

For data centre operators to fulfil their promise of power continuity to tenants and avoid costly disruptions, facilities need to be supported with reliable redundancy equipment. In traditional 2N redundancy models, power systems are duplicated and run simultaneously, so that if the primary power source fails due to unplanned maintenance or unexpected faults the redundant system can take over without interrupting the critical load.

However, this approach involves investing in twice as much electrical equipment – generators, inverters, UPS and switches – despite it not always being actively used, driving high capital and operational expenditure, as well as space inefficiencies. Importantly, this model of simply adding more redundancy infrastructure to guarantee resilience has become unsustainable in the AI age.

RETHINKING RESILIENCE

Reassessing how data centres ensure resilience can help operators make significant progress in reducing PUE while



adapting to the demands of AI workloads. Catcher architecture offers an alternative approach to resilience that doesn't require duplicating every element of power infrastructure.

During a power fault, the load transfers from the normal path to the redundant path, which 'catches' it online, without any interruption or voltage drops. Although this design requires additional hardware in the form of static transfer switches (STS) to be placed between the UPS and the load, Socomec estimates that catcher architecture uses 30 per cent less equipment and 42 per cent fewer batteries overall.

Catcher architecture can be paired with more efficient UPS technology to further reduce the energy overhead of resilience. While many existing units operate at about 95 per cent efficiency, modern UPS systems can reach up to 99.1 per cent and can even support grid interaction – automatically drawing more power when voltage is high to improve overall facility efficiency. And since less energy is lost through the UPS overall, operators can make better use of the power available to support growing AI workloads.

'Rather than designing power protection around a fixed level of demand from the outset, modular UPS architectures allow extra capacity to be introduced as requirements develop.'

ROOM FOR MANOEUVRE

Space is already at a premium in today's data centres, particularly in existing facilities that were not designed around today's levels of compute density. If large amounts of floorspace are taken up by redundancy equipment, operators have less room for the IT capacity needed to support AI, limiting their facilities' revenue potential.

This means that UPS design needs to be considered not only in terms of efficiency but also physical footprint and ease of installation.

More compact systems can deliver greater power capacity within a smaller area, while reducing installation requirements such as rear or side clearance can make that capacity easier to accommodate in constrained environments. Design choices that simplify cooling and maintenance can also free-up surrounding space that needs to be reserved for a UPS. For operators looking to increase capacity within an existing facility, opting for a compact UPS design can help them prioritise capability over hardware.

'UPS systems must demonstrate flexibility by smoothly integrating with existing power infrastructure. Compatibility with different battery technologies and existing electrical systems gives operators more freedom to adapt their approach as the needs of the facility change.'



SCALING POWER PROTECTION

Operators shouldn't assume that the power requirements of AI will stay static. As deployments develop, the facility's needs may change, making flexibility an integral part of its resilience. Rather than designing power protection around a fixed level of demand from the outset, modular UPS architectures allow extra capacity to be introduced as requirements develop. This can help operators respond to changing demand without replacing or substantially redesigning the wider system each time more capacity is needed.

Modularity can also strengthen resilience by making parts of the UPS easier to replace without compromising the wider system. Hot swappable modules allow operators to replace a faulty component without taking the entire UPS offline, so that it continues to provide protection as problems are being fixed.

Finally, UPS systems must demonstrate flexibility by smoothly integrating with existing power infrastructure. Compatibility with different battery technologies and existing electrical systems gives operators more freedom to adapt their approach as the needs of the facility change.





SUPPORTING GROWTH

As AI places greater demands on data centre infrastructure, operators must rise to the challenge by considering how their resilience architectures can support the wider performance of the facility, beyond simply guaranteeing uninterrupted supply. Modern UPS systems help operators make better use of available space and power, while providing the flexibility needed to safely accommodate the fluctuating demand patterns of AI. The result is a resilience strategy that leaves more of the data centre's resources available for the IT capacity and which generates business value rather than merely acting as back-up. ■



COLIN DEAN

Colin Dean is managing director at Socomec. With over 30 years' experience in the electrical and data centre industry, he leads the UK and Ireland business, supporting the delivery of resilient power infrastructure for critical environments.

Colt becomes first European based provider to integrate last mile connectivity with Equinix Fabric Marketplace

Colt Technology Services has become the first European based provider to integrate its last mile connectivity services with the Equinix Fabric Marketplace. Fabric Marketplace gives enterprises a single destination to quickly discover, learn about and connect with service providers that can enhance their distributed networking deployments for multicloud and AI workloads.



The enhanced collaboration combines Colt's extensive global network footprint with Equinix Fabric, enabling businesses and enterprises to establish high performance, secure connectivity from third-party locations, including office buildings and data centres, directly into Equinix facilities. Designed to eliminate traditional procurement and deployment complexities, the integration gives enterprises faster access to the scalable network infrastructure required to support cloud first strategies, data intensive workloads and AI-driven growth.

Bringing Colt's last mile connectivity services into the Fabric Marketplace simplifies how organisations across Europe and Asia discover, procure and deploy connectivity services. Through the Fabric Marketplace experience, customers can instantly identify available Colt services, access pricing information, request quotes and submit order requests.

Vantage and Heron Power Electronics announce GW scale 800VDC agreement to accelerate AI infrastructure delivery

Vantage and Heron Power Electronics have announced a strategic, GW scale manufacturing capacity agreement. Under the agreement, Vantage has secured over 1GW of manufacturing capacity of Heron Link, Heron's solid-state transformer (SST), to deploy in its hyperscale data centre campuses. The collaboration is designed to give AI and cloud customers a faster, more efficient and scalable way to deploy the high density infrastructure required for next generation computing.



As AI rack densities approach MW class power requirements and the grid needs additional flexible and interactive loads, customers face growing constraints related to power availability, electrical equipment, deployment timelines and cost. Heron Link addresses those challenges by converting medium voltage grid power directly to 800VDC with integrated energy storage, reducing the number of conversion stages and simplifying the path from the grid to the rack.

Schneider Electric, AVEVA, Nscale and Start Campus bring the NVIDIA Omniverse DSX Blueprint into a live AI factory environment



Schneider Electric, AVEVA, Nscale and Start Campus have announced the successful pilot of an end-to-end digital twin developed using the NVIDIA Omniverse DSX Blueprint, part of NVIDIA DSX Sim. The Immersive Twin pilot marks a significant step toward bringing Schneider Electric's digital twin technology, powered by AVEVA software, into real world AI infrastructure, allowing Nscale to accelerate token production and optimise token efficiency.

The deployment is focused on Nscale's operations within the SIN01 data centre on the 1.2GW SINES Data Campus in Portugal, developed by Start Campus. As one of Europe's largest and most sustainable AI factories, the SINES Data Campus is designed to support successive generations of AI infrastructure.

It will soon become home to more than 66,000 NVIDIA Rubin graphics processing units (GPUs), representing one of the largest AI deployments in the region. By combining engineering and operational data with advanced power and cooling simulations and workload requirements, the Schneider Electric and AVEVA Immersive Twin provides a unified view of data centre performance to help identify opportunities to optimise infrastructure before changes are implemented on-site.

PROJECTS & CONTRACTS IN BRIEF

Salute, a member of the NVIDIA Partner Network (NPN), has secured customers to support more than 20GW of AI infrastructure capacity globally, following rapid growth in its work supporting the deployment and operation of high density NVIDIA accelerated computing infrastructure for some of the world's leading AI providers.

Shropshire Council has agreed a new five year contract with MLL Telecom, with the option to extend for a further two years. The scope of the new contract includes a refresh as well as further expansion of the network to 87 sites, encompassing all Shropshire and partner locations.

Epsilon Telecommunications and Megaport have partnered to expand Epsilon's network footprint in North America. Customers can now access Megaport's compute, network and storage across the US and Canada via Epsilon, with software defined connectivity to key data centre, cloud and internet exchange hubs across the region.

ServiceNow and the Bundesliga have expanded their partnership, with ServiceNow becoming the league's Official CRM Partner, following early results that have significantly improved how it handles fan enquiries.

Schneider Electric



Schneider Electric's RapidScale Modular Data Center Reference Design 115 is a new reference design that provides a validated, pre-engineered blueprint for high density AI infrastructure, streamlining on-site logistics and planning while accelerating deployment times for AI factories globally.

The design brings together the NVIDIA Vera Rubin platform and innovative liquid cooling technology, built to maximise processing power while keeping energy use under control. By using advanced fluid cooling with smart temperature controls, it ensures data centres provide peak compute performance without risking overheating or wasting power.

By integrating modularity across the broader infrastructure architecture, this approach transitions operators toward repeatable, risk tested AI infrastructure blueprints rather than relying solely on custom engineered data centres. The blueprints combine power, liquid cooling, IT, software and lifecycle services into a scalable, industrialised delivery model.

CLICK HERE to find out more.

www.se.com

Siemon



Siemon's LightStack and LightStack 8 ultra-high density fibre plug and play system delivers superior density, port access and cable management in a sleek, modern platform. This makes it ideal for today's advanced data centre or enterprise networking environments.

The black LightStack enclosure and connectivity range is the Base-12 format, while the grey LightStack 8 enclosure and connectivity range is the Base-8 format. Both configurations support singlemode and multimode fibre.

Siemon's innovative solution supports up to 144 fibres per rack unit in a 1U enclosure and enables easy access and management of connectivity from the front or rear. It offers various mounting options and high capacity management clips that are designed to support the maximum capacity of cables as well as simple access. It also includes a range of plug and play transition modules and pass through adaptor plates designed to deliver a level of performance exceeding today's standards.

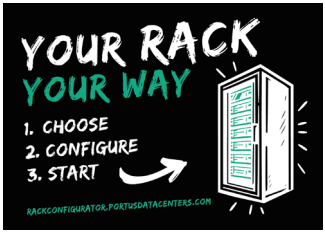
To find out more **CLICK HERE.**

www.siemon.com

Portus Data Centers

Rack configuration made simple

The Portus Data Centers interactive digital rack configurator allows colocation customers greater freedom and transparency when planning rack requirements. The dedicated portal enables them to independently assemble rack solutions quickly and easily, compare different options and submit sales inquiries online.



In German or English, customers can select the appropriate Portus Data Centers location, rack widths and units, combining them into a customised solution. Available standard capacities are clearly displayed.

Power supply, pre-wiring, patch panels, cable management and other components can be configured via an interactive 3D view. The rack unit matrix enables precise component placement on the front and back of the rack. Technical catalogue rules automatically validate the configuration, while utilisation, set-up costs and monthly expenses are displayed transparently.

The configurator is particularly well suited for small to medium sized enterprises (SMEs), start-ups, IT service providers, system integrators and managed service providers, while also offering a simple entry point when planning infrastructure for AI inference workloads.

CLICK HERE for more information about the Portus Data Centers interactive digital rack configurator.

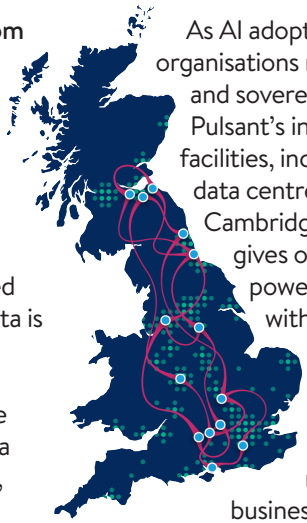
www.portusdatacenters.com

Pulsant

Powering the UK's AI future from the edge

AI is driving a fundamental shift in digital infrastructure demand. While training models often happens in hyperscale environments, the rapid growth of AI inference is creating a need for compute closer to where data is generated and put to use.

Pulsant is uniquely positioned to support this shift. Its UK-wide platform, bringing together data centres, connectivity and cloud, provides the low latency, high performance digital infrastructure businesses need to deploy AI applications closer to users, devices and critical operations.



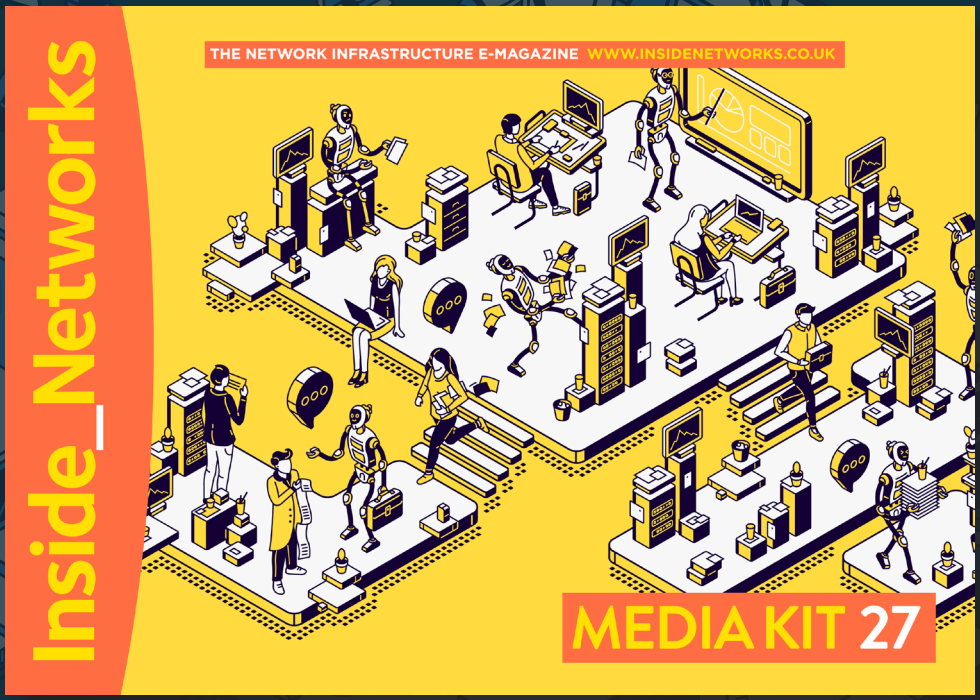
As AI adoption accelerates, organisations need flexibility, resilience and sovereignty over their data. Pulsant's investment in high density facilities, including its Milton Keynes data centre within the Oxford-Cambridge innovation corridor, gives organisations access to powerful digital infrastructure without relying solely on heavily constrained London locations.

By extending enterprise grade compute into the regions, Pulsant is helping businesses unlock the full potential of AI while keeping data secure, connected and close to where value is created.

For more information **CLICK HERE.**

www.pulsant.com

All you need to know



CLICK ON THE COVER TO VIEW THE 2027 MEDIA KIT

Rand and deliver

Oskar Lampe of BCS Consultancy explains why South Africa is becoming a key hub for digital infrastructure and what this means for Africa's growing data centre market



'Africa has long been viewed as one of digital infrastructure's most promising long-term opportunities. What's changing is not the scale of that opportunity, but how quickly it is becoming reality.'

For the last decade, the global data centre conversation has focused on familiar markets. North America has continued to expand, Europe has grown beyond its traditional hubs into regions such as the Nordics, Italy and Iberia and, more recently, the Middle East has emerged as a destination for hyperscale investment, supported by government strategies and significant energy resources.

WHY SOUTH AFRICA?

Africa has long been viewed as one of digital infrastructure's most promising long-term opportunities. What's changing is not the scale of that opportunity, but how quickly it is becoming reality. As investment accelerates and the market matures, developers, operators and technology companies are increasingly looking at Africa as a region for growth

‘Over the past decade, South Africa has attracted the largest concentration of hyperscale and colocation investment on the continent, supported by established cloud regions, a mature professional services sector and a supply chain that is becoming increasingly experienced in delivering mission critical infrastructure.’



today, not sometime in the future.

While growth is taking place across multiple African markets, South Africa has emerged as the continent’s most established digital infrastructure hub. It’s a natural starting point for organisations looking to build a long-term presence in Africa. However, its position has not happened by chance.

Over the past decade, South Africa has attracted the largest concentration of hyperscale and colocation investment on the continent, supported by established cloud regions, a mature professional services sector and a supply chain that is becoming increasingly experienced in delivering mission critical infrastructure. This combination provides a level of market maturity that is difficult to find elsewhere in Africa, while still offering significant room for growth.

SIX APPEAL

The market is expected to more than double in value over the next six years, increasing from approximately $2.55bn in 2025 to $5.28bn by 2031. Cloud adoption is accelerating alongside digital transformation across both the public and private sectors, while rising mobile data consumption, the continued growth of fintech and increasing demand for data sovereignty are creating sustained demand for secure, low latency digital infrastructure.

As digital infrastructure investment gathers significant pace in countries such as Kenya, Nigeria and Morocco, many international developers, operators and technology companies are looking for a regional base that combines established capability with access to emerging markets. South Africa is well placed to fulfil that role, providing both a strong domestic market and a platform from which organisations can support growth from across the wider continent.



capacity creates the foundation for further digital transformation across the continent.

This cycle of investment is changing how international organisations view the region. Africa is no longer being assessed solely for its future potential but increasingly for the opportunities already emerging today.

NEXT GENERATION

Across more established data centre markets, AI is beginning to influence the type of infrastructure being built, with higher rack densities, increased power requirements and more advanced cooling strategies. These requirements are rapidly becoming part of mainstream data centre design rather than specialist considerations.

Africa has an opportunity to incorporate many of those lessons from the outset. Much of the infrastructure needed to support future AI adoption is still in the planning or development phase, allowing developers to design for future workloads rather than retrofit existing facilities. That creates the potential to build infrastructure that is not only capable of supporting today's demand but is also prepared for the rapid evolution of AI over the coming years.

MORE CONNECTED

The pace of investment is also being shaped by improving connectivity. New subsea cable systems, together with continued expansion of terrestrial optical fibre networks, are strengthening connections both within Africa and between Africa and the rest of the world. As latency reduces and resilience improves, the continent becomes increasingly attractive for cloud providers, enterprise applications and digital businesses.

Those investments are reinforcing one another. Better connectivity enables greater cloud adoption, which increases demand for local data centre capacity. As more organisations move critical workloads into the cloud, the need for resilient infrastructure located closer to end users continues to grow. In turn, that additional

LEARNING CURVE

Rapid growth inevitably brings delivery challenges. Power, specialist skills and

global supply chains will all influence how quickly projects move from concept to operation, just as they have across Europe, North America and the Middle East.

Power remains one of the defining considerations for new developments and Africa is not immune to those challenges. Although South Africa's electricity situation has improved considerably in recent years, demand for digital infrastructure is increasing rapidly and connection strategies will continue to influence site selection and commercial viability. Like other mature markets, power is becoming a strategic part of project planning rather than simply a utility requirement.

The same is true of people. South Africa has a strong community of engineering, commercial and project management professionals, but specialist mission critical experience is still developing because the sector itself remains relatively young. Investing in local capability will therefore become just as important as investing in physical infrastructure. Supply chains present another consideration. Many critical components including transformers, switchgear, generators and cooling systems remain dependent on global manufacturing and logistics.

BENEFIT CHECK

None of these issues are unique to Africa. The advantage is that the market has the opportunity to benefit from lessons already learned elsewhere, adopting proven approaches while developing delivery models that reflect local conditions and support long-term, sustainable growth.

As international investment accelerates, success in Africa will depend on more than replicating delivery models developed elsewhere. Understanding local procurement practices, regulatory environments, supply chains and delivery capability will become increasingly important as projects grow in scale and complexity. South Africa provides a strong platform from which to build that understanding.

TAKING THE ADVANTAGE

As the market matures, competitive advantage will come not simply from global experience but from combining it with genuine local knowledge, trusted partnerships and a long-term commitment to the region. Organisations that invest early in building those capabilities will be best placed to support Africa's next phase of digital infrastructure growth. ■



OSKAR LAMPE

Oskar Lampe is regional director at BCS Consultancy. Originally from South Africa, he is leading BCS' African expansion, spearheading the establishment of a new regional office and helping to build upon the company's long-term strategy.

WE HOPE YOU HAVE ENJOYED

Inside_Networks

COMING UP IN THE
JAN/FEB 27 ISSUE:
SPECIAL FEATURES

> COPPER CABLING SYSTEMS
> CABLE MANAGEMENT AND LABELLING

TO FIND OUT MORE ABOUT PROMOTION WITHIN THESE FEATURES CLICK HERE

> ALL THE LATEST NEWS, VIEWS, COMMENT AND ANALYSIS

> 2026 REVIEWED AND A LOOK AHEAD AT WHAT'S IN STORE FOR 2027

> THE HIDDEN COST OF POOR LABELLING AND WHY SMALL DETAILS CAN HAVE LONG-TERM CONSEQUENCES

> HOW COPPER CONNECTIVITY TECHNOLOGIES HAVE EVOLVED TO MEET CHANGING DATA CENTRE DEMANDS, FROM TRADITIONAL STRUCTURED CABLING TO HIGH SPEED COPPER INTERCONNECTS

> WHY SPEED TO POWER IS CHANGING DATA CENTRE DEVELOPMENT

> THE WAITING GAME FOR DATA CENTRE CAPACITY AND HOW BUSINESSES CAN BEAT IT

> HOW SMART CABLE MANAGEMENT CAN HELP FUTURE PROOF DATA CENTRES

> THE EVOLUTION OF AI CLUSTER INTERCONNECT DESIGN AND THE MULTIPATH RELIABLE CONNECTION (MRC) APPROACH

> HOW AI IS FORCING DATA CENTRE OPERATORS TO RETHINK POWER

> MOVES, ADDS AND CHANGES IN THE CHANNEL

> NETWORK INFRASTRUCTURE CASE STUDIES FROM AROUND THE WORLD

> THE LATEST PRODUCT, SYSTEM AND SERVICE DEVELOPMENTS

FOR A FREE SUBSCRIPTION CLICK HERE