DCS
DATACENTRE SOLUTIONS

SPECIAL EDITION OSLO 2026

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

HTTPS://DATACENTRE.SOLUTIONS

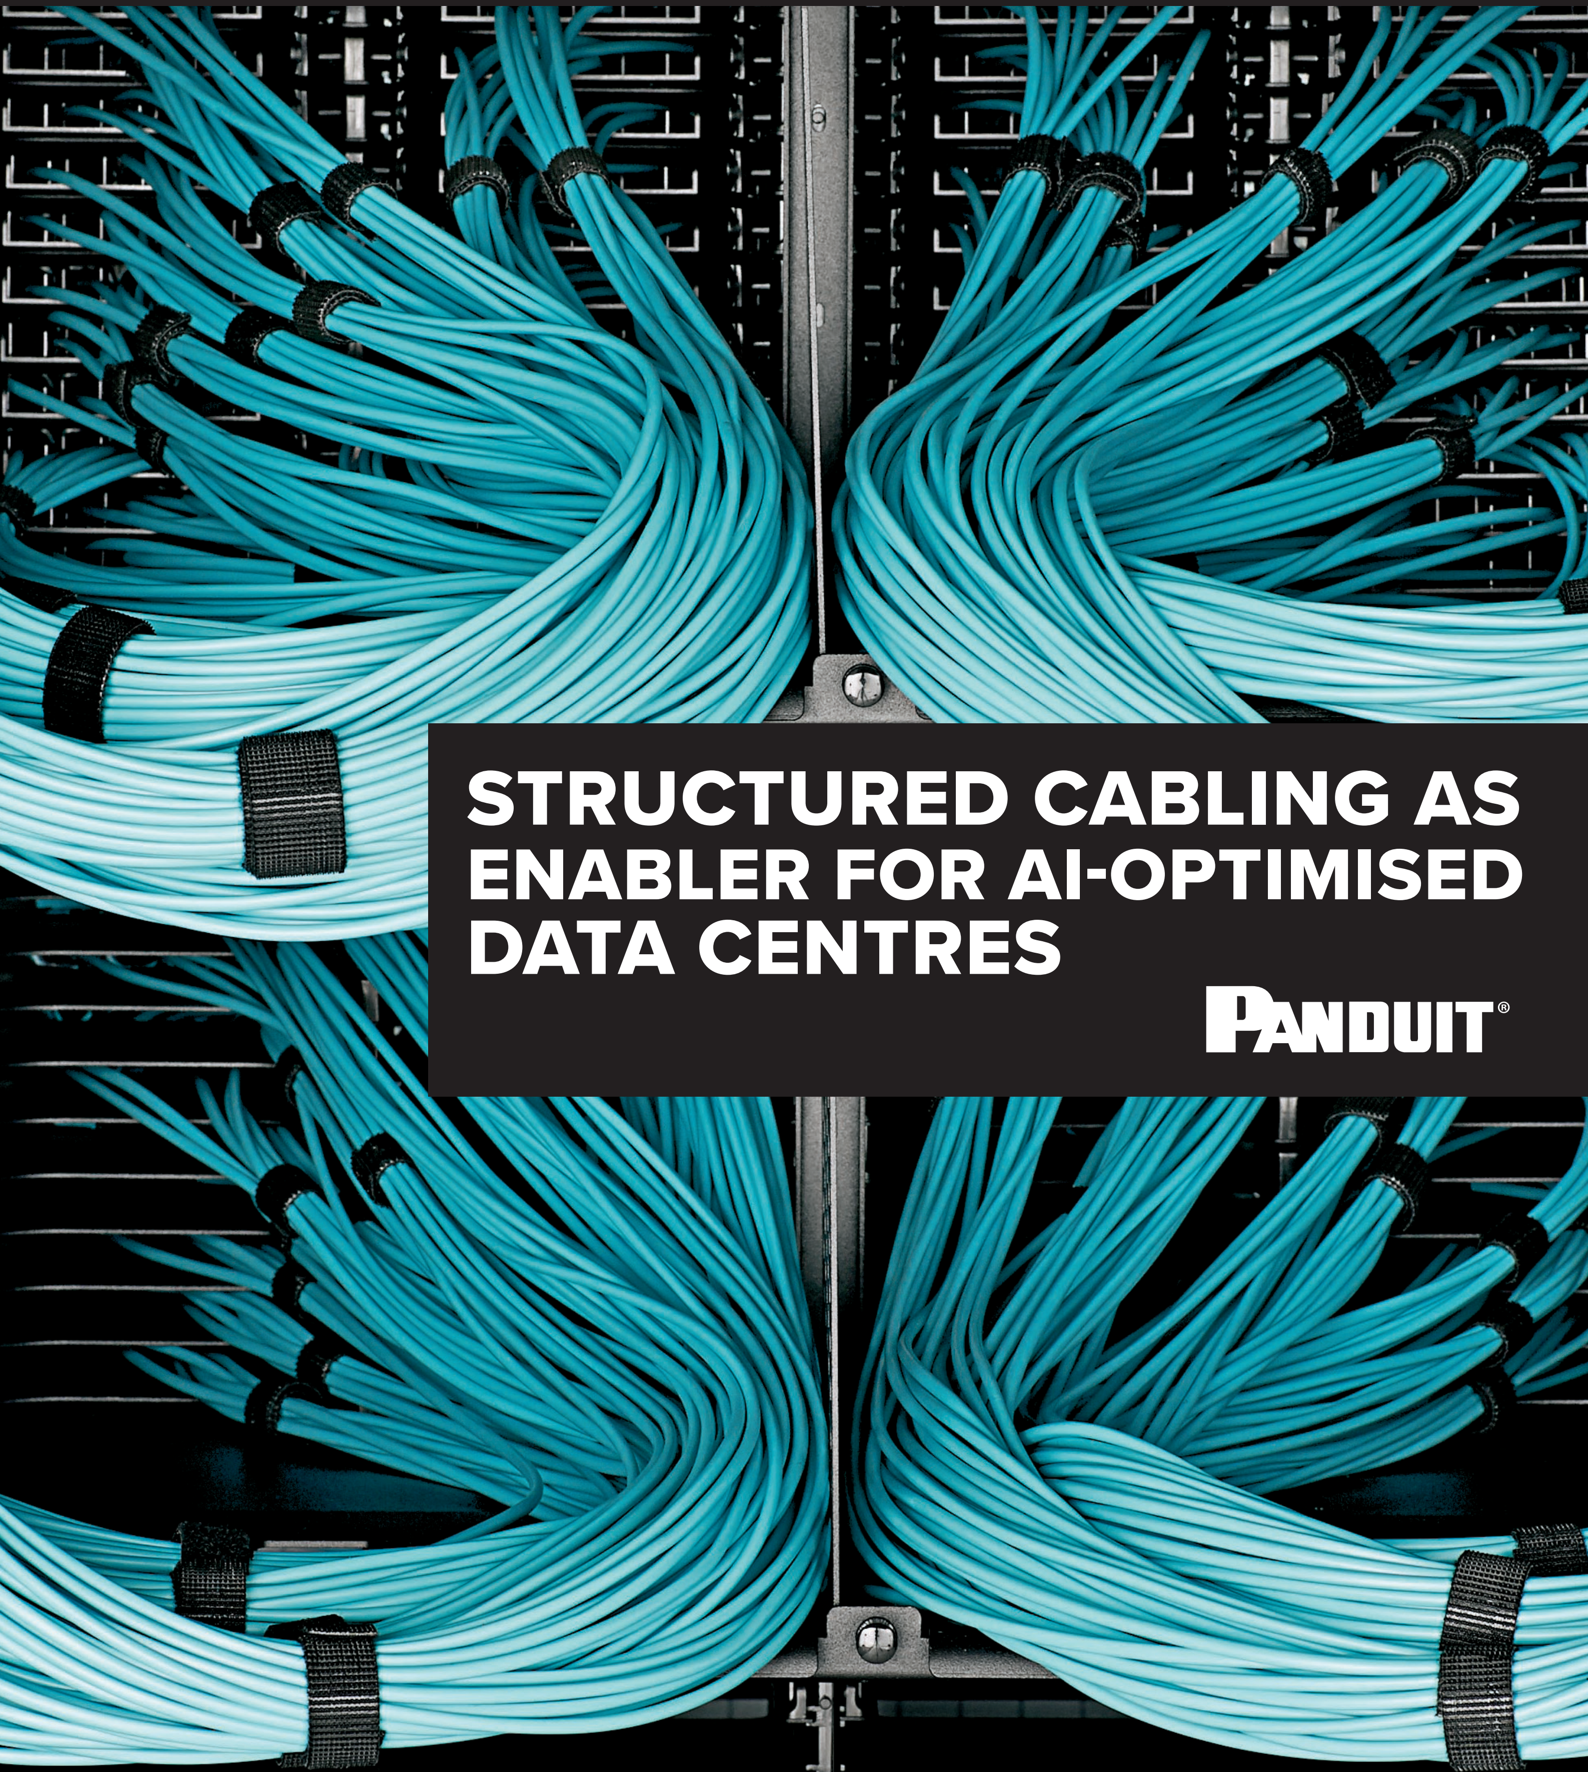

STRUCTURED CABLING AS ENABLER FOR AI-OPTIMISED DATA CENTRES

HPC SERIES

Cool, Smart, Reliable

TUDERTECHNICA®

www.tudertechnica.com

VIEWPOINT

BY PHIL ALSOP
EDITOR

AI – time to transition from LLMs to inferencing?

FOR THE PAST TWO YEARS, the data centre conversation has been dominated by large language models and the immense hyperscale campuses required to train them. Training clusters built to support organisations such as OpenAI have reshaped expectations around rack density, liquid cooling and multi-hundred-megawatt grid connections. But training is only the opening chapter. The next act - and arguably the more commercially pervasive one - will be inference.

As AI adoption matures, the balance of activity will shift from building ever larger models to running them, repeatedly and at scale, closer to users and machines. That transition has significant implications for the size, geography and design philosophy of data centres.

Inference workloads are persistent, latency-sensitive and widely distributed. Autonomous systems, smart manufacturing lines, healthcare diagnostics, retail analytics and connected infrastructure do not have the tolerance for round-trip delays to distant hyperscale regions. As edge AI becomes embedded in everyday operations, compute will need to sit nearer to population centres, logistics hubs and industrial clusters.

This does not spell the end of the hyperscale campus. Training frontier models will still require vast, power-dense facilities in locations with abundant energy and land. However, the gravitational centre of growth may broaden. Instead of a small number of enormous builds in select regions, we are likely to see a proliferation of mid-sized, highly efficient edge facilities — tens rather than hundreds of megawatts — positioned to serve metropolitan and regional demand.

In markets such as the United Kingdom and across the Europe, that could mean renewed focus on secondary cities. Locations previously overlooked in favour of London, Paris, Frankfurt or Dublin may become attractive for inference hubs, particularly where fibre connectivity, renewable energy access and supportive planning frameworks align. Proximity

to users, not just proximity to cheap power, will shape decision-making.

Design priorities will also evolve. Edge AI facilities must balance performance with operational simplicity. They will need modular scalability, advanced cooling for GPU-accelerated workloads, and robust physical and cyber security — but often within tighter footprints and closer to communities. Noise, visual impact and grid integration will face greater scrutiny.

Crucially, the sustainability debate will intensify. Distributing compute more widely risks increasing total energy consumption if not carefully managed. Yet inference can also be more energy-efficient per transaction than repeatedly querying distant centralised models. Intelligent workload placement — deciding what runs locally and what returns to core regions — will become a strategic discipline.

The AI era is entering a decentralised phase. Those who can blend hyperscale strength with agile, edge-ready design will be best placed. As inferencing moves from novelty to norm, the map of digital infrastructure may look very different indeed.

Contents

50

26

DCS DATACENTRE SOLUTIONS

Editor
Philip Alsop
+44 (0)7786 084559
philip.alsop@angelbc.com

Senior B2B Event & Media Executive
Mark Hinds
+44 (0)2476 718971
mark.hinds@angelbc.com

Director of Logistics
Sharon Cowley
+44 (0)1923 690200
sharon.cowley@angelbc.com

Design & Production Manager
Mitch Gaynor
+44 (0)1923 690214
mitch.gaynor@angelbc.com

Graphic Design & Multimedia Assistant
Harvey Watkins
harvey.watkins@angelbc.com

Publisher
Jackie Cannon
+44 (0)1923 690215
jackie.cannon@angelbc.com

Circulation & Subscriptions
+44 (0)1923 690214
circ@angelbc.com

Directors
Scott Adams: CTO
Sukhi Bhadal: CEO

Published by: Angel Business Communications Ltd, 6 Bow Court, Burnsall Road, Coventry CV5 6SP. UK
T: +44 (0)2476 718970 E: info@angelbc.com

Angel BUSINESS COMMUNICATIONS

Datacentre Solutions is published 10 times a year on a controlled circulation basis in Europe, Middle East and Africa only. Subscription rates on request.
ISSN 2756-1143 (Online)

From efficiency to value:

Why Power Compute Effectiveness (PCE) must lead data center decision-making metrics in a power-constrained era

For more than a decade, the data center industry has relied on a familiar set of performance metrics to guide design, operation, and sustainability narratives.

BY PAUL QUIGLEY, PRESIDENT AND CHIEF STRATEGIC RELATIONS OFFICER, AIRSYS COOLING TECHNOLOGIES

THESE METRICS emerged during a period of relative abundance – when power was accessible, land was available, and growth was constrained primarily by capital and engineering execution. But that world no longer exists – and, in truth, it may never have existed in the way the industry assumed.

Across Europe and North America, data center development is detrimentally constrained by power availability, grid interconnection delays, permitting red-tape, water awareness, noise concerns, and growing community resistance. In this environment, efficiency alone is no longer sufficient.

The industry needs to, and is being forced to, move beyond measuring how efficiently power is consumed and begin measuring how effectively power is converted into economic and societal value. The distinction between efficiency and effectiveness is not semantic, it is foundational. This article makes a deliberate and necessary argument: data center metrics must be reordered based on the decisions they are intended to support. In a power-constrained world, that order must be:

- Power Compute Effectiveness (PCE) must precede
- Infrastructure Utilization Efficiency (IUE), which must precede
- Power Usage Effectiveness (PUE).

This is not a rejection of existing metrics. It is an acknowledgment that their relevance depends on context. And simply put, that context has changed.

Metrics exist to enable decisions

Metrics are not academic exercises. They exist to answer specific questions and enable specific decisions. When metrics are applied outside the context they were designed for, they obscure reality rather than clarify it.

To understand why the industry must rethink its metric hierarchy, it is essential to define clearly and without ambiguity, what each metric actually tells us, and what it does not.

Power Usage Effectiveness (PUE): Optimizing the building

What question PUE answers:
- How efficiently does the facility deliver electrical power to IT equipment?"

What PUE reveals:
- Energy lost to things like cooling, power conversion, and facility overhead.
- Mechanical and electrical efficiency of the building
- Operational practices and behaviors

What decisions PUE informs:
- Cooling system selection
- Electrical distribution design
- Facility retrofits to reduce operating costs.

What PUE does not reveal:
- Whether IT equipment is producing useful work
- Whether available power is fully utilized
- Whether the site has economic or strategic value

End result:
PUE is an engineering efficiency metric and an important one. It helps operators reduce waste and lower operating expense. It does not determine whether a data center is valuable, expandable, or even viable in a constrained market. A facility can exhibit an excellent PUE and still be fundamentally constrained by power availability or thermal ceilings.

Infrastructure Utilization Efficiency (IUE): Optimizing IT assets

As compute density increased and workloads became more dynamic, the industry correctly recognized that facility efficiency alone was insufficient. This led to growing interest in utilization-based metrics.

What question IUE answers:
"How much of the installed IT capacity is actually performing useful work?"

What IUE reveals:
- Idle or underutilized servers
- Capital efficiency inside the data hall
- Opportunities for workload optimization

What decisions IUE informs:
- IT upgrade or update strategies
- Server consolidation initiatives
- Workload strategic logistical planning

What IUE does not reveal:
- Whether additional power can be delivered to the site
- Whether existing power is stranded by non-IT constraints
- Whether the facility can grow within regulatory or community limits

End result:
IUE is an IT capital efficiency metric, and an excellent one. It improves return on deployed hardware but remains blind to the most significant constraint facing modern data centers: access to usable power. Utilization matters… but only after power exists.

Power Compute Effectiveness (PCE): Revealing Value

The defining constraint of today's data center market is no longer space, capital, or cooling efficiency. It is power availability.
This reality demands a metric that answers a different question entirely.

What question PCE answers:
"How much useful compute is delivered for each unit of scarce electrical power?"

What PCE reveals:
- Whether power is fully monetized or stranded
- How much compute potential is suppressed by non-power constraints
- The true productive yield of a site's electrical allocation

What decisions PCE informs:
- Whether a brownfield site is worth upgrading or abandoning
- Where capital should be deployed to unlock stranded power
- Which sites should be prioritized in constrained markets
- Whether permitting, community engagement, or infrastructure upgrades will generate real returns

End result:
PCE is a value-revealing metric. It connects power, compute, and economics in a single outcome-oriented lens.

This is where PCE fundamentally differs from PUE and IUE.

The only metric that correlates to ROIP

Return on Invested Power (ROIP) has emerged as a critical financial concept in a world where megawatts. not square meters, are the scarce resource.

> Across Europe and North America, data center development is detrimentally constrained by power availability, grid interconnection delays, permitting red-tape, water awareness, noise concerns, and growing community resistance. In this environment, efficiency alone is no longer sufficient

ROIP asks a simple but unforgiving question: "What return do I generate from each megawatt I secure?"

Only one metric answers the technical half of that equation.
- PUE reduces cost per megawatt-hour
- IUE improves use of installed IT capital
- PCE increases value per megawatt

PCE is therefore the only metric that correlates directly and structurally to ROIP. It determines whether power – once contracted, permitted, and delivered – can actually be converted into revenue-generating compute.

Without PCE, ROIP remains an abstract financial aspiration. With PCE, it becomes measurable and actionable.

Why stranded data centers exist

Across both the United States and Europe, a significant number of data centers possess "contracted" or "deliverable" power that cannot be fully utilized. These sites are often described as "stranded or power-limited," yet the limitation is rarely the grid connection itself.

Instead, power is stranded by:
- Thermal operating assumptions
- Conservative inlet temperature envelopes
- Old building decisions that were optimized for past legacy workloads
- Design constraints unrelated to actual chip and silicon capability

PUE and IUE frequently, in fact almost always mask this reality. A stranded facility can appear efficient and reasonably utilized while remaining economically stagnant.

PCE exposes the truth.
It reveals how much power exists, how much compute should be possible, and how much value is being left unrealized.

Sustainability, permitting, and social license

Modern data center opposition rarely centers on efficiency metrics. Communities and regulators care about:
- Power consumption
- Water usage
- Noise
- Heat rejection
- Grid stability

These concerns are not solved by small incremental improvements in PUE or utilization. They are addressed by maximizing compute delivered per unit of societal impact.

Improving PCE reduces:
- Cooling intensity per unit of compute

METRIC	METRIC EVALUATES	METRIC OPTIMISES	IDEAL USERS	ENABLES DECISIONS IN	CANNOT REVEAL
PUE Power Use Effectiveness	How efficiently a facility delivers power to IT equipment	Facility Energy & Operating Cost	Facility Engineers Operators	Cooling Architecture Selection Electrical Design OPEX Reduction Efficiency Retrofits	*IT Productivity* *Power Scarcity* *Stranded Capacity* *Site Value*
IUE Infrastructure Utilisation Efficiency	The amount of installed IT capacity doing useful work	IT Capital Utilisation	IT Leadership Operations Teams	Server Consolidation Workload Placement Refresh Cycles CAPEX Improvements	*Power Availability* *Permitting Limits* *Site Expansion Potential*
PCE Power Compute Effectiveness	The amount of useful compute delivered per unit of scarce power	Power-to-Compute Value	Executives Investors Utilities Policymakers	Brownfield vs. Greenfield Investment Stranded Power Recovery Site Prioritisation CAPEX Deployment	*Operational or Financial Value Creation Under Power-Strained Conditions*
ROIP Return On Invested Power	The return generated per megawatt secured	Financial Return on Power	Boards Investors Utilities	Capital Allocation Power Contracting Strategy Long-term Infrastructure Planning	*Technical Causes of Underperformance without PCE Input*

EFFICIENCY METRICS OPTIMISE SYSTEMS THAT ALREADY EXIST
EFFECTIVENESS METRICS DETERMINE WHETHER SYSTEMS ARE WORTH EXISTING AT ALL

- Water consumption per delivered workload
- Mechanical scale and acoustic footprint
- The need for incremental power infrastructure

In this way, PCE becomes the bridge metric, connecting engineering decisions to permitting outcomes and community acceptance.

Why this conversation has been delayed

The industry has been slow to adopt PCE for a simple reason: PCE removes comforting abstractions. This distinction may be uncomfortable, but it directly exposes the differences between efficiency and effectiveness. Efficiency metrics optimize systems that already exist;

Effectiveness metrics determine whether systems are worth existing at all.

PCE strips away the layers and forces uncomfortable but necessary questions:

- Why does power exist here that cannot be monetized?
- Which constraints are physical, and which are inherited assumptions?
- Where is value being suppressed rather than unlocked?

These questions are no longer optional.

Reordering the conversation

The future of digital infrastructure will be shaped not by how efficiently power is consumed, but by how effectively power is converted into compute that serves economies and societies.

That requires a clear hierarchy:

- PCE — Determines value and viability
- IUE — Optimizes IT capital within available power
- PUE — Optimizes facility efficiency within utilized infrastructure

Applied in this order, metrics illuminate reality rather than obscure it.

Conclusion

The data center industry is entering an era defined by power scarcity, regulatory friction, and social accountability. Metrics designed for an age of abundance must evolve accordingly. Power Compute Effectiveness does not replace existing metrics. It precedes them. It answers the question that must be resolved before any other optimization matters: can available power be converted into meaningful compute and economic value?

Until that question is answered, efficiency – by itself - is irrelevant.

Keeping the environment under control

Forbo Flooring Systems is an international market leader with a wealth of experience in ESD flooring offering a range of ESD solutions which can serve the highest demands and requirements of controlled environments.

REDUCING the generation of electrostatic charges is the main purpose of control measures in ESD protected areas (EPA). In areas such as data centres, the right floor covering plays a crucial role. It not only drains electrostatic charges from personnel and equipment, but it also reduces the generation of charges where they occur, at the interface between the soles of shoes and the floor.

For this segment, we recommend our Colorex SD I EC ranges of static control flooring. There are conductive and static dissipative options along with a Plus collection that is ideal where downtime or subfloor repairs are an issue.

Colorex SD/ EC tiles

With Colorex SD & EC, static charges flow easily through the dense network of tiny conductive veins that run through the whole thickness of the tile.

The charge is transmitted via the conductive adhesive and securely discharged to earth via the copper strip. It's a completely natural system that needs no volatile chemical anti-static additives to aid conductivity.

The conductivity is not affected by changes in temperature or humidity.

It can be installed by fully adhering it directly to a subfloor or bonding it to a raised access floor panel. It can also be laid as part of a loose lay tile system.

Colorex is available with three levels of conductivity - High performance, Performance and Standard. The loose lay slip-resistant version (R11 Plus) completes the comprehensive Colorex high-tech solutions offer.

The benefits of Colorex:

- The technology ensures a permanently conductive floor covering guaranteeing optimal performance throughout it's life, regardless of the humidity level of the area. Colorex fully complies with all ESD standards.
- The surface of Colorex can be fully restored and repaired extending it's useful life.
- The low plasticizer content in Colorex tiles creates excellent dimensional stability and prevents shrinkage.
- The low plasticizer content also ensures extremely low emissions and outgassing making Colorex the right choice for cleanrooms. Colorex fully complies with all cleanroom standards, confirmed by the Fraunhofer institute.
- The dense construction of Colorex and the Colorex plus loose lay system provides excellent resistance to heavy loads making it an ideal solution for commercial and industrial environments.
- Colorex is available in tile format, a prerequisite for raised access floors found in many commercial environments such as data centres and server rooms.
- There is an option to add safety signs and guiding with Colorex signal and Colorex signal glow.
- Being one of the most trusted brands on the market, you can have peace in mind that Colorex complies to every standard and norm that is required today.

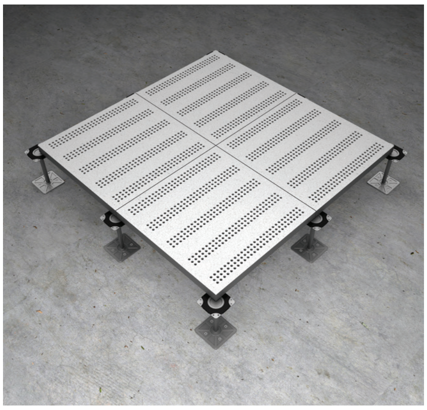

➤ An advanced technical flooring system specifically designed to control static discharge in sensitive areas such as data centres. Colorex is not only an advanced technical solution, it is also aesthetically pleasing, enhancing any interior.

Creating better environments

We constantly strive to produce sustainable flooring systems that create better environments everywhere. At the same time, we help take care of the natural environment through our commitment to sustainable development, responsible raw material procurement and manufacturing processes.

Our Colorex collections are manufactured using 100% green electricity and are REACH compliant.

Of all the electricity we buy, 100% comes from renewable sources.

This means all our vinyl production sites, including our Colorex plant, are part of an effective environmental management system and achieve ISO14001 certification.

One year on: How Elevate is redefining data centre infrastructure at speed

Where did that year go? It feels like yesterday that Elevate - Future Faster launched at Data Centre World 2025

SINCE then, we've been busy working with operators, partners and integrators to understand exactly where white space designs strain under growth, how airflow, power and patching need to adapt, and how deployment can be accelerated without adding operational risk. Now we're back for year two at Data Centre World London (ExCeL), 4–5 March, Stand B180, not with "new for the sake of new", but with solutions that close the gap between what modern data centres demand and deliverable infrastructure.

Elevate was created to make density practical, control intuitive and scale repeatable. The guiding principle is simple: more density, more control, more scale - without complexity creeping in through the back door. That's why the platform brings together fibre, racks, aisle containment, power and security as a cohesive ecosystem rather than a catalogue of parts. In our second year, that ecosystem has grown to offer wider choices for highdensity fibre, stronger airflow strategies, smarter power and access control - and all of it engineered to install cleanly, label clearly and operate predictably.

How Elevate addresses today's DC challenges

Modern enterprise and data centre environments share the same reality: higher density, faster change cycles and tighter operational guardrails. Elevate is designed to help operators keep pace, technically and operationally.

- **Densification without chaos.** Port counts continue to climb while real estate stays flat. Operators need patch fields, ODFs and racks that accommodate more in the same or smaller footprint - without turning cable routing and access into a tangle. Elevate's high-density connectivity - spanning VSFF, MPO-based architectures and highcapacity ODFs - delivers port density while maintaining front-of-rack access, sensible bend radius and clear, consistent labelling. The goal isn't simply "fit more"; "it's fit more and manage better."
- **Thermal containment that actually works.** As loads rise, so do heat and variability. Adhoc airflow tactics quickly reach their limits. Elevate's hot and cold aisle containment options are engineered to integrate with rack layouts, door and roof interfaces, brush entries and cable ways, keeping supply and return paths distinct so the thermal model stays stable.
- **Power visibility and control at scale.** Energising a rack isn't enough anymore. Teams need to monitor usage, balance phases and, when required, switch safely under change windows. Elevate's high-density intelligent power options provide the outlet counts modern builds demand, with the visibility and control that make daytwo operations more predictable. It's about risk reduction through data - and designs that anticipate growth.
- **Deployment speed and predictability.** Live programmes can't stall because field splicing overruns, or a routing compromise triggers rework. Elevate prioritises preconnectorised fibre and engineered pathways so builds move quickly and consistently. Pre-terminated trunks and pre-populated ODF trays reduce time on site, cut variability, and increase first-time-right outcomes - especially important for large rollouts and repeatable pods.
- **Physical security and auditability.** As more stakeholders touch infrastructure, who accessed what, when and why becomes central to governance. Elevate's intelligent rack locking integrates at rack level to deliver rolebased access and event logging, improving both security posture and audit readiness, without slowing legitimate work.
- **Operational clarity as you scale.** Growth often fails not at design, but during day-two operations when documentation, labelling and manageability drift. Structured labelling, clear patch presentation and tray-level guidance are engineered into the Elevate platform to support consistent work practices as estates expand. The objective is a platform that scales without accumulating ambiguity.
- **Fast logistics and availability.** Lead time and delivery confidence can make or break a deployment. Elevate's approach is to maintain fast, reliable availability across core configurations - with clear build standards and predictable options - so planning isn't disrupted by supply uncertainty. The result: design intent stays intact, and rollout schedules remain credible.

Advancing the Elevate platform for 2026

- **VSFF ultra-high-density pre-connectorised fibre optics.** As patching fields become denser, VSFF connectors offer high port density in familiar rack unit footprints up to 3456 fibres in 1U while preserving front access ergonomics and labelling.

Preconnectorised assemblies take the speed and predictability further, minimising on-site splicing and reducing test overheads. The practical win is a cleaner, faster build that remains serviceable at scale.

- **Hot aisle containment.** For facilities optimised around hot air capture and reuse, Elevate's hot aisle solutions are designed to create consistent separation with doors, roofs and interfaces that integrate to rack, cable and power routes. The outcome is a more stable thermal profile and higher cooling efficiency, especially as loads rise and hardware mixes evolve.

- **High-density intelligent power management.** The power layer now needs to be instrumented, not just provisioned. Elevate's highdensity intelligent power options focus on outlet availability, metering and, where appropriate, switching, so teams can balance, monitor and automate reliably. The value is fewer surprises during change windows, supported by meaningful telemetry.
- **Intelligent rack locking solutions.** Physical access is a security control that must scale with estates. Elevate integrates rack-level access control to support role-based permissions and auditable events, aligning with compliance requirements while keeping hands-on work flowing efficiently.
- **High-density ODF with pre-connectorised trays and cables.** Patching fields grow, but the requirements remain the same: clarity, accessibility and repeatability. Elevate's ODF architecture emphasises tray-level structure, managed bend radius and clear labelling, while pre-connectorised options shorten installation time and reduce post-install rework.

Alongside these 2026 additions, we're also showcasing the enhanced features of the Elevate DCR Rack Series, cold aisle containment, and MPO high-density pre-connectorised solutions. The DCR Series now brings refinements that make higher-density builds easier to construct and maintain - enhanced airflow features, better cable routing and improved load handling - while cold aisle options support retrofit scenarios and mixed environments with predictable performance.

MPO-based architectures continue to excel for backbone and spineleaf

designs, helping standardise on repeatable, fast pathways.

None of these are "checkbox" features. Each addresses a specific operational pressure - density, airflow, power control, access security or deployment speed - so operators can design once, repeat confidently and maintain clarity as they scale.

Practical advantages operators experience with Elevate

Operators tell us the wins are immediate and cumulative. Install windows shrink because pre-connectorised options reduce on-site splicing and test cycles. Patch fields stay legible at high density thanks to tray design, front access ergonomics and consistent labelling. Thermal behaviour becomes predictable when containment is engineered to the rack and cable interfaces rather than added as an afterthought. Power events are less dramatic because monitoring and switching are planned in from the start. Access audits move faster when rack-level controls integrate with operational processes. And across it all, the fact that fibre, racks, containment, power and security are designed to work together reduces the friction that often appears when mixing components from multiple sources.

These are not theoretical benefits. They play out in the small decisions, week after week: a patch that routes cleanly the first time; a cabinet that cools as modelled; a change window that closes on schedule because the right telemetry was available; a compliance check that passes smoothly because the access trail is credible. Each result is modest on its own; together, they form a platform that helps teams go faster without losing control.

Where Elevate meets real-world requirements - DCW London

To truly understand infrastructure, you need to see how it's built, how it routes, and how it performs. If you're at Data Centre World London (ExCeL), visit Stand B180 to work through the details with the Elevate team. Open ODF trays and see how identification and bend radius are managed at density.

Trace routing paths in DCR racks and explore airflow accessories. Walk the aisle containment interfaces and understand how the mechanical joins support a consistent thermal model. Review intelligent power options for outlet density, metering and switching, and test the physical access experience with intelligent rack locking. Come and discuss your upcoming projects with the team and explore how Elevate's ecosystem - fibre, racks, aisle, power and security - can support your growth, densification or refresh objectives in 2026.

This is why more enterprise and data centre operators are turning to Elevate: it's a platform engineered to make higher density environments manageable, to accelerate deployment without sacrificing quality, and to maintain operational clarity as estates evolve.

More density. More control. More scale.

See it on B180.
Year two for Elevate isn't about arriving; it's about accelerating - closing the gap between what modern data centres demand and what infrastructure can deliver, at speed and with confidence. If your 2026 roadmap includes growth, densification or a data centre refresh, make time for Stand B180. See the VSFF, ODF, containment, power and security layers working together. Meet the team, pressure-test the details, and decide where Elevate can simplify your next build or upgrade.

And because a visit should be useful and enjoyable: enter Elevate's on-stand competition for a chance to win a pair of Apple AirPods while you're there.

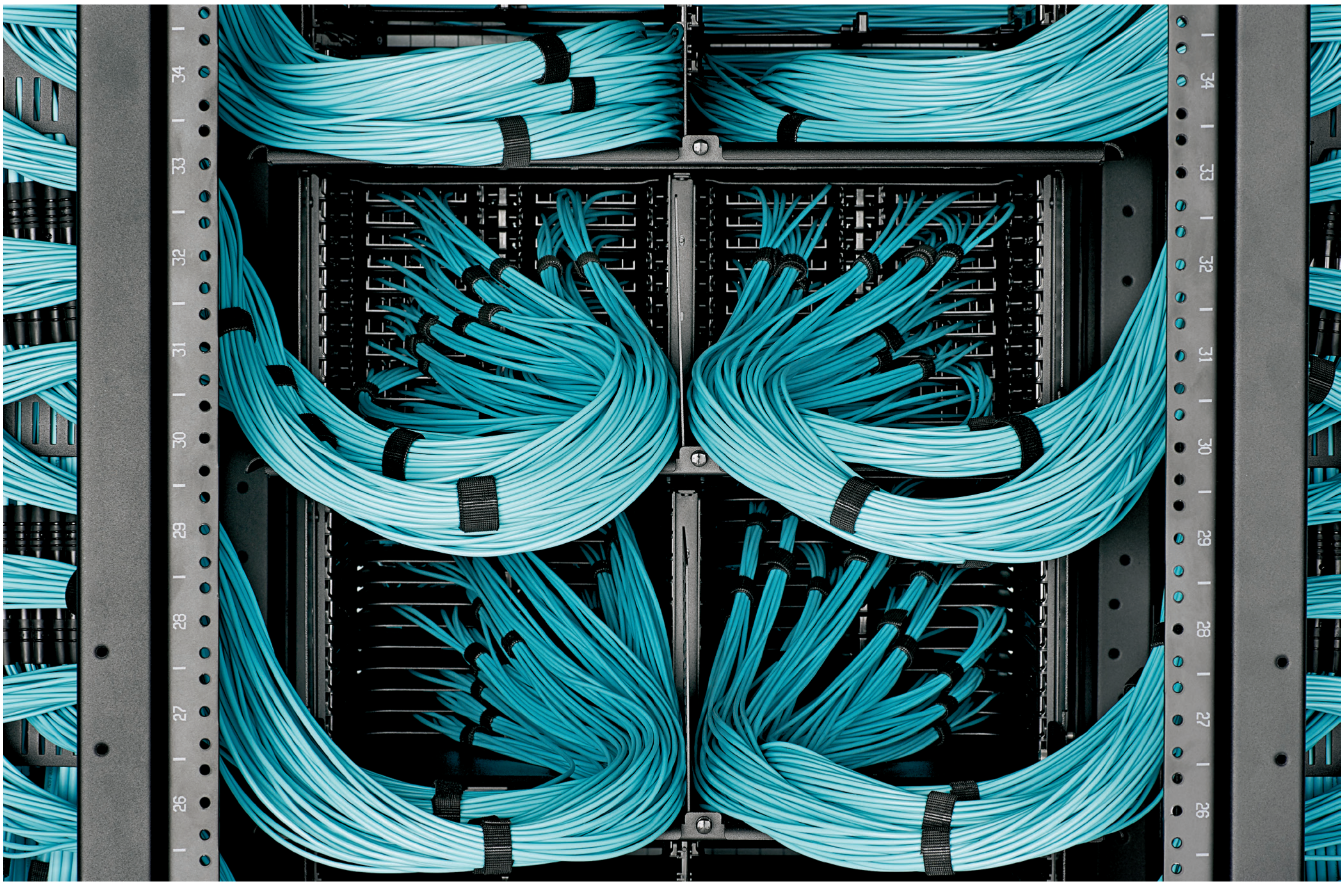

Structured cabling as enabler for AI-optimised data centres

AI workloads are not just reshaping compute, they are redefining the physical foundations of the data centre. As training and inference fabrics push network speeds from 400 Gb/s to 800 Gb/s and forward to 1.6Tb/s, the tolerance for inefficiency in the underlying infrastructure is rapidly disappearing

BY MICHAEL AKINLA, PANDUIT

WHAT WAS ONCE considered a passive layer, structured cabling, has become a critical enabler of performance, scalability, and operational control. In this environment, deterministic design, optical integrity, and modular deployment models are essential to supporting high-density GPU clusters and evolving network architectures.

Structured cabling, when properly engineered and validated, is emerging as a key strategy for meeting these demands while future-proofing data centre investments.

The move to 800G Ethernet in AI back-end networks is driving a shift from duplex LC to parallel-optic MPO connectivity. Current GPU servers and leaf switching platforms commonly present MPO-08 interfaces operating as eight duplex lanes (16 fibres) to deliver 800G aggregate throughput under IEEE 802.3df, using either 800GBASE-SR8 over multimode fibre or 800GBASE-DR8 over single-mode fibre.

Relative to LC duplex architectures, this increases fibre count per link by at least 4×. When multiplied across a pod-scale deployment containing large numbers of GPU servers, total installed fibre density typically climbs toward ~8× that of a conventional data centre network. This is a direct consequence of lane-based optics and is not optional if bandwidth scaling is to be maintained. Panduit's Base-8 structured cabling is aligned to this lane geometry, enabling consistent fibre mapping, simplified management, and a clean uplift path from 400G to 800G without forcing a backbone redesign.

Fibre density is exploding and unmanaged direct connect topologies don't scale

The rapid adoption of 400Gb/s and 800Gb/s network speeds requires significantly more fibre links. AI clusters rely on APC multi-fibre MPO connections for server to leaf links, and your more traditional single-mode MPO connections for leaf to spine links, which means fibre volumes have increased exponentially. Without a structured approach, data centres risk excessive cable congestion, increasing difficulty in maintenance and reducing airflow optimisation.

Latency concerns are largely a myth in AI SuperPods

AI back-end networks are typically deployed as tightly clustered SuperPods with short optical reaches, often <50m. Propagation delay across this distance is well below 250 ns, so any incremental length in a structured topology is negligible compared with the dominant latency situations: switch serialization, buffering, and FEC decode latencies. Panduit's measurements confirm that structured cabling adds optical interfaces but does not create meaningful latency overhead at SuperPod distances, and that controlling slack in structured pathways can actually reduce avoidable propagation length that accumulates in direct-connect builds.

Structured cabling adds discrete connector interfaces, so the relevant engineering question is optical loss budget rather than latency. IEEE channel models explicitly allocate connectivity loss headroom, approximately 1.5 dB for multimode channels and ~2.5 dB for single-mode channels in standards-compliant designs.

Panduit's testing with NVIDIA 800G SR8 and DR8 transceivers shows that when connector losses are maintained within these allocations and standard cleanliness and inspection regimes are followed, structured cabling links remain within Bit Error Rate (BER) and power-margin limits required by IEEE 802.3df.

The result is straightforward: structured cabling is channel-safe for AI workloads when engineered to spec, with performance governed by loss control, connector quality, and disciplined maintenance, not by topology choice.

Implementing structured cabling for AI workloads

To maximise ROI and ensure reliable operation through multiple speed generations, structured cabling for AI fabrics should be designed around five practical requirements:

- **Scalability and modularity**
  A modular cross-connect architecture using high-density MPO trunks and patching enables deterministic scaling without re-pulling backbone fibres. This approach supports iterative pod growth and incremental speed adoption while maintaining stable fibre routing and polarity schemes.
- **Optimised pathways and fibre management**
  High fibre density is a mechanical and thermal issue as much as a connectivity issue. Structured cabling consolidates multiple links into high-count trunks and terminates into patch fields, sharply reducing pathway fill, improving cable-tray serviceability, and preserving airflow. Panduit frames offer structured consolidation as a primary control for preventing congestion at AI scale.
- **Reliability and reduced downtime**
  Structured systems enforce bend-radius protection, separate fixed trunks from equipment cords, and anchor slack in controlled storage zones. When combined with rigorous labelling and documentation in line with TIA practices, this produces traceable circuits, lowers handling-related fault rates, and reduces MTTR by speeding physical-layer isolation.
- **Future-ready, high-density connectivity**
  Using 16-fibre/8-lane MPO channels for 800G establishes a standards-aligned physical baseline compatible

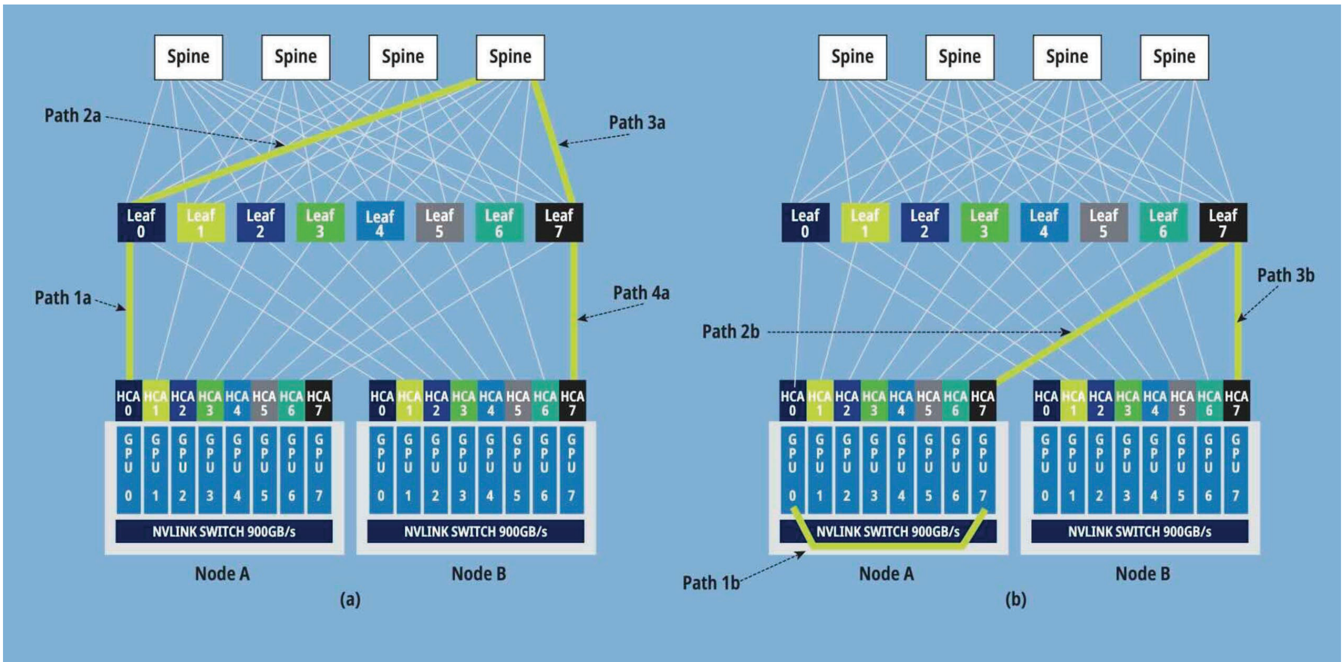

Ilustrates a rail-optimized topology configuration. Yellow traces show the signal patch to compare latency when using (a). The Leaf and Spine switches and (b) The NVLINK and Leaf switches.

➤ Figure 1. Rail Optimised Topology Configuration.

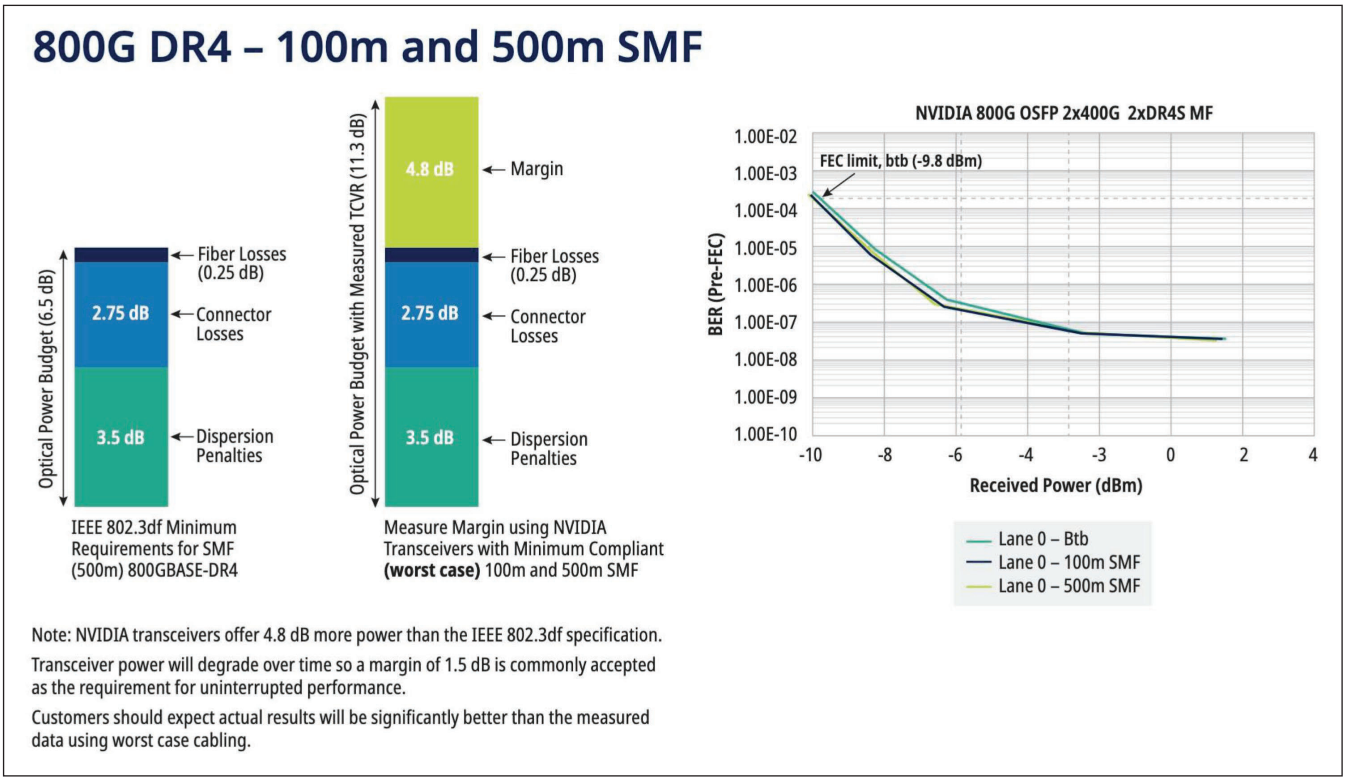

➤ Figure 2. 800G DR4 structured-cabling power budget and BER headroom (Panduit/NVIDIA). Measured power-budget breakdown and BER curves demonstrating structured-cabling compliance margin under DR8 optics.

with 400G while enabling a clean uplift path to higher lane rates. Base-8 structured topologies preserve this forward compatibility and minimise stranded assets during generational upgrades.

- **Energy efficiency and sustainability** Because AI SuperPods are short-reach environments, multimode fibre is often a viable physical layer and supports lower-power multimode optics (up to 15% lower) compared with equivalent single-mode transceivers. Panduit positions structured, reusable fibre backbones as part of long-term efficiency strategy by reducing recabling churn and enabling more controlled lifecycle upgrades.

Structured cabling vs. point-to-point cabling in AI fabrics

Point-to-point cabling can appear simpler at very small scale, but in AI fabrics it creates an increasingly unpredictable physical layer. As pods grow, unmanaged direct runs introduce large slack volumes, uncontrolled bend conditions, inconsistent routing, and high pathway occupation, all of which degrade serviceability and raise handling-induced failure probability. Expansion is also mechanically disruptive because each uplift adds new discrete circuits, compounding congestion and limiting airflow. Over the system lifecycle, these factors translate directly into higher operational cost driven by longer troubleshooting cycles and a greater tendency toward recabling during moves and upgrades.

Structured cabling avoids these failure modes by imposing a repeatable physical topology with modular cross-connects, controlled slack and routing, and predictable upgrade mechanics. In the AI context, where scaling is expected, not exceptional, this deterministic physical-layer model is the lower-risk engineering choice.

The road to 1.6 Tb/s and beyond

The data centre's physical layer must evolve alongside next-generation AI switching and optics. As 1.6 Tb/s standards mature, co-packaged optics (CPO) will shorten electrical paths and intensify the need for clean, high-density fibre distribution that can be re-terminated without large-scale backbone replacement. At the same time, emerging compact connector systems such as MDC-class formats will increase front-panel port counts, making structured polarity control and routing discipline even more critical. Higher-order modulation such as PAM-4 will further extend throughput over existing fibre plants, reinforcing the value of a structured backbone that can be reused through multiple generations rather than repeatedly replaced.

Across these trends, the core requirement remains stable: AI data centres need a physical-layer architecture that is scalable, loss-controlled, mechanically serviceable, and standards-compatible.

AI fabrics are pushing data-centre physical layers into a regime of very high bandwidth, extreme fibre density, and rapid lifecycle scaling. The transition from LC duplex to MPO parallel optics, the ~8× increase in fibre volume within AI pods, and the move toward 1.6 Tb/s make structured cabling fundamental to performance and operability. Panduit/NVIDIA validation shows that, within IEEE loss budgets and with disciplined installation practice, structured cabling preserves BER margin and introduces no meaningful latency overhead in SuperPod-class networks.

Talent scarcity, trust, and insider risk: Why secure hiring must become a strategic capability for data centres

The data centre sector is undergoing one of its most rapid growth phases ever, driven by cloud expansion, AI tools, and increasing business digitisation. This surge in demand for capacity and new centres brings a major challenge: finding enough skilled, trustworthy people.

BY ASHLEY FERNANDES, EXECUTIVE DIRECTOR, CORPORATE SECURITY SERVICES, ICTS UK & IRELAND

COMPETITION for talent is intensifying across engineering, operations, security, and facilities management. Experienced professionals have many options, career paths are less straightforward, and hiring processes are lengthening. Organisations often rush to fill positions, which, while understandable, can inadvertently increase insider risk.

Insider threat is often seen as malicious behaviour, but in reality, it frequently results from mistakes, poor judgement, policy breaches, or individuals under personal or professional stress. In data centres, where availability, integrity, and confidentiality matter most, small errors can rapidly escalate into big problems.

Although the sector has invested heavily in physical security, perimeter controls, and advanced electronic measures, these alone do not address risks stemming from hiring unsuitable people or inadequate assessment processes.

Risk begins long before access control systems come into play if the wrong individuals are hired. Therefore, secure hiring must be treated as a distinct security discipline, not merely an administrative task.

In today's tight labour market, technical skills often dominate hiring decisions. However, a more sustainable approach prioritises character, judgement, and values over qualifications. Increasingly, operators seek to understand not only candidates' credentials but also how they behave under pressure, handle responsibilities, and whether their attitudes align with company culture.

This requires more than standard background checks. Comprehensive identity verification, employment and education history validation, right-to-work confirmation, and criminal record screening should be standard. For high-risk positions, deeper vetting, structured behavioural interviews, and scenario-based assessments can better gauge suitability.

Importantly, insider risk evolves over time. Changes in a person's circumstances, financial issues, burnout, disengagement, or external pressures can alter behaviour. Hence, many organisations adopt lifecycle assurance models including probation monitoring, periodic re-screening for critical roles, and confidential channels for reporting concerns. Technology is also aiding this shift. Identity-centric access systems, workforce analytics, and behavioural monitoring create smarter, proportionate trust strategies. When used responsibly, these tools help detect early warning signs and enable timely interventions before problems escalate.

At the executive level, the link between labour shortages and insider risk is now a strategic concern. The focus is not just on how many staff are needed, but whether organisations consistently recruit, select, and retain trustworthy people. This is where industry groups can play a vital role. Organisations like the Data Centre Alliance (DCA), by uniting operators, service partners, and subject matter experts, can promote shared values and practical guidance on secure recruitment and people-focused security.

While data centres are built with concrete, steel, and cutting-edge technology, they rely on people for operation and maintenance. Organisations investing in trust-based hiring and thorough vetting will be better positioned to reduce insider risk and build safer, more resilient, and sustainable data centre operations in the future.

The data centre industry in 2026 – repatriation, resilience, and regional rebalancing

The climate has been tough for businesses, with slow growth, high costs, cyber risks and geopolitical uncertainties all contributing to the challenge. More than ever, businesses must innovate to survive and grow, and digital infrastructure has a key role.

ROB COUPLAND, CEO OF PULSANT

THE PAST YEAR has been pivotal for cloud strategy, with repatriation gaining momentum due to shifting legislative, geopolitical, and technological pressures. This trend continues, with a growing focus on data sovereignty. As we settle into 2026, the stage is set for it to be the year of repatriation, resilience, and regional rebalancing.

Repatriation goes mainstream

Data repatriation remains a key trend, with businesses re-evaluating cloud adoption due to rising costs, regulatory pressures, and concerns over data sovereignty. Many are moving workloads from public cloud to private cloud, on-prem, or colocation to gain better control, cost efficiency, and compliance. While cloud remains essential, a hybrid model combining public, private, and on-prem solutions is emerging as the future of enterprise IT strategy.

For UK businesses, this has prompted many to shift workloads from global hyperscalers to domestic providers, creating hybrid infrastructure blends, especially as data sovereignty becomes a top priority. Latest research indicates 87% plan to repatriate some or all of their workloads over the next two years. This year businesses will focus even more on visibility, data locality, sovereignty, and transparency. However, managing a hybrid or mixed infrastructure poses challenges, as the era of static infrastructure fades away.

The continued impact of cyber security breaches

Several recent high-profile cybersecurity breaches in UK organisations have led to increased awareness of where data is stored and how it's managed, particularly within the supply chain. Many businesses still don't fully understand where their data resides, how it's processed, or how it's backed up. These breaches have renewed the focus on resilience, particularly in terms of the speed of recovery after an incident, rather than just prevention.

This is now influencing data centre policies and services, with disaster recovery and backup becoming more standard. UK data centre providers are already stepping up with more transparent, compliant platforms to meet growing customer demands for visibility and secure infrastructure.

We'll also see the true impact of the Cyber Security & Resilience Bill during 2026, as its guidelines start to shape industry practices and expectations. While the Bill introduces stronger cyber security measures and clearer regulations, it also brings increased reporting standards and red tape that could prove burdensome. Policy developments will favour well-prepared, security-focused operators.

The AI hype cycle hits reality

AI-driven data centre demand surged in 2025, especially in hyperscale facilities, as interest in AI technologies reached new heights.

However, as the hype settles, businesses are starting to evaluate real-world AI uses and determine what digital infrastructure is truly needed to support their AI goals. This also brings inference AI and sovereign AI into the picture, further complicating the landscape, with Edge computing emerging as a key beneficiary.

Although hyperscale demand will no doubt continue, demand for specialised, inference-optimised storage platforms will become more significant.

Regional rebalancing & the rise of Edge

While London and the South East will likely maintain their dominance in the UK data centre market, there's growing interest in regional locations, spurred by initiatives like Tech Towns and AI Growth Zones, which encourage investment outside the capital. Taking the economic benefits of technology clusters into the regions is something we champion.

We'll see increased availability of Edge data centres near UK metros, opening up new opportunities for sectors like smart manufacturing and transport. These regional edge facilities will offer more sustainable, cost-effective infrastructure, contributing to a more balanced national digital economy.

We're going through a period of significant change, with rising interest in the data centre sector as it plays an increasingly critical role in supporting industry and society.

Providers who prioritise transparency, regional diversification, and realistic AI enablement will be best positioned to lead the way.

AI and the enhanced fire risk within modern data centres

The introduction of larger capacity server racks to meet growing Artificial Intelligence demands has significantly increased fire risk in data centres.

BY DAVE SMITH, CBS BUSINESS DEVELOPMENT MANAGER FOR UK & EUROPE

EARLY DETECTION and rapid activation of automatic extinguishing or suppression systems are critical to preventing facility downtime, business interruption and consequential losses.

Understanding the Standards

BS 5839-1:2025 provides the foundation for fire detection and alarm system design in non-domestic premises. However, where areas of increased fire risk exist, such as Electronic Data Processing (EDP) installations like data centres, BS 6266 becomes the essential reference standard.

BS 5839-1 specifies maximum coverage of approximately 100m^2 for optical smoke point detectors, based on a 10.6m x 10.6m grid with devices positioned no more than 7.5m from walls or each other. However, BS 6266 recognises the high air flow environments typical of EDP areas and requires significantly closer spacing, reducing detector coverage to just 25m^2-a 4:1 ratio requiring substantially higher detection density.

This closer spacing uses a grid with detection points 2.5 metres from side walls and 5 metre centres. Fire system designers working with BS 6266's Table I.1 will be familiar with additional mitigating factors that may reduce coverage further still.

The case for aspirating detection

Data centres classified as High or Critical risk under BS 6266-where equipment is high-value or purpose-built, operations are not easily transferable, data requires continual remote backup, and business interruption carries serious consequences-demand more comprehensive detection strategies. This is where Aspirating Detection Systems (ASD) excel.

ASD systems actively draw air through a pipe network with sampling holes or capillaries to a central detector that monitors for minute smoke particles. EN54-20 defines ASD performance through three sensitivity classifications:

- **Class A - Very High Sensitivity**: Detects fire at its earliest stages in environments where air handling heavily dilutes smoke. Essential for computer rooms, data centres and clean rooms.
- **Class B - Enhanced Sensitivity:** Provides earlier warning than standard detectors for protecting high-value, vulnerable or critical equipment in areas where smoke is diluted by airflow or building volume.
- **Class C - Normal Sensitivity:** Matches standard point smoke detection sensitivity, primarily used to overcome installation and application challenges.

Implementation strategy

ASD systems should be installed on the Return Air of room air handling-known as Primary Detection, typically at Class A sensitivity. This captures smoke following its most likely travel path in data halls with managed airflow. Secondary Detection using ASD should also be considered for ceiling installations within room spaces and ceiling/floor voids, mimicking standard point detection at Class B or Class C sensitivity.

Conclusion

Understanding risk is paramount. Enhanced detection using Aspirating Smoke Detection is essential, supplemented by appropriate extinguishing or fire suppression systems for complete protection. As cooling systems, airflow management and power requirements continue evolving with AI demands, fire system designers must leverage this information to deliver effective solutions to emerging fire protection challenges.

The 2AM Troubleshooting Call

That Could Have Been Prevented

Outdated cable documentation turns routine maintenance into all-hands emergencies because nobody knows what is plugged in where. ***Rapid*ID™ eliminates the guessing game:** pre-labelled patch cords with barcode scanning deliver real-time accuracy.

RAPIDID™

Reduce Time and Cost
of cable documentation compared to traditional manual approaches

Automate Labour-Intensive and Error-Prone Cable Documentation Process
that can lead to an outage

Free Up Network Engineers
to focus on more rewarding and strategic initiatives

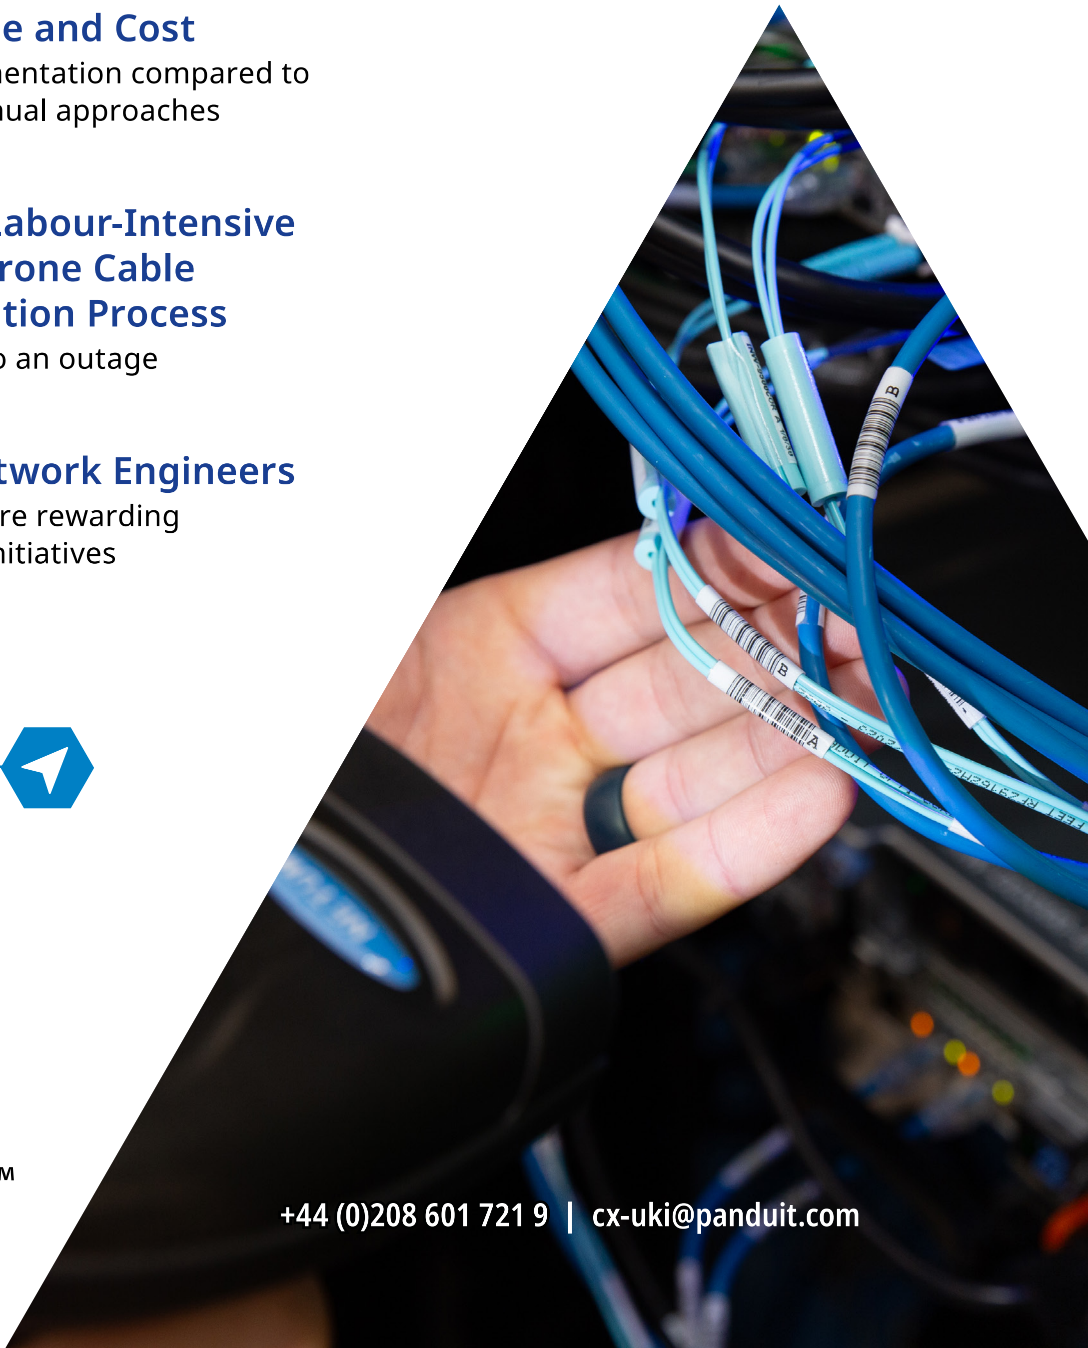

For more information, visit
www.panduit/rapidid

PANDUIT™

+44 (0)208 601 721 9 | cx-uki@panduit.com

The end of guesswork: Engineering precision for the AI factory

The paradigm for designing data centers is undergoing a fundamental transformation, driven by the unprecedented demands of artificial intelligence (AI).

BY SHERMAN IKEMOTO, SALES DEVELOPMENT GROUP DIRECTOR AT CADENCE DESIGN SYSTEMS

THE traditional approach – fragmented design cycles, siloed engineering disciplines, and reliance on outdated rules of thumb—is no longer viable. Today's AI factories are not merely buildings housing servers; they are highly integrated, complex systems where power, cooling, and IT infrastructure are deeply interdependent. The reason for this evolution is the significant increase in power consumption driven by the demands of modern AI workloads, which require more energy-intensive hardware and advanced cooling solutions to maintain efficiency and performance. This new reality requires a shift from approximation to precision, a transition that our digital twin technology is engineered to enable.

For decades, datacenter design has relied on safety margins at multiple points to compensate for design uncertainty—uncertainty created by a fragmented design process and the lack of advanced tools capable of delivering greater precision. This practice, while once a necessary precaution, now represents a grave inefficiency.

As rack densities escalate from 20-kW to over 100-kW to support next-generation GPUs, the financial and operational costs of overprovisioning become unsustainable. Guesswork is a liability the industry can no longer afford. The path forward lies in simulation-driven design, where validated, physics-accurate digital models replace assumptions with data-backed certainty.

This is the strategic imperative behind digital twins. By creating behaviorally accurate models of critical AI infrastructure, such as the latest high-performance GPUs and AI SuperPOD, designers and operators can virtually test, validate, and optimize their facilities before a single piece of hardware is installed. These are not static drawings but dynamic multiphysics models augmented with AI surrogates, that can rapidly simulate power consumption, thermal behavior, and airflow under real-world workloads with unprecedented precision. This allows for design directly to service-level agreements (SLAs), minimizing waste and maximizing performance, with confidence.

Cadence's collaboration with GPU vendors extends this principle from the chip to the entire data center. Just as we have partnered for decades to design the world's most advanced silicon, we are now extending that collaboration to the system and facility level. The integration with NVIDIA Omniverse further accelerates this "extreme co-design" approach, enabling multidisciplinary teams to work concurrently within a shared, simulation environment.

The opportunity doesn't stop at design. It is imperative to also optimize operations and lifecycle management. A digital twin that validates a facility's initial design is a platform that can also simulate maintenance, test upgrades, and manage/optimize performance and efficiency in real time. The inability to reliably and accurately predict gaps in data center performance, and determine their root cause, is the single greatest bottleneck in deploying AI factories at scale. By embedding physics-based simulation accelerated with AI inferencing into the design and operational lifecycle, we are closing those gaps. We are replacing guesswork with the science-backed optimization from the Cadence Reality Digital Twin Platform, ensuring that the infrastructure powering the future of AI is built on a foundation of precision.


AIRSYS
PowerOne™
Precision Cooling for the AI Era.
A New Line of Innovative Cooling Solutions Designed to Unlock Stranded Capacity.
PowerOne™ cooling systems were developed to help data centers extract maximum value from available power. By reducing cooling energy overhead and supporting high-density deployments, PowerOne improves Power Compute Effectiveness (PCE), ensuring a greater share of incoming power reaches the rack.
This is more than just ROI — it's a Return on Invested Power.
Improve Your PCE. Increase Your ROIP.
Explore PowerOne at www.air-sys.uk

Aalberts: the lifeblood of every modern data centre

In the complex data centre supplier landscape, which as we all know is experiencing unprecedented growth and with new suppliers entering the DC market every day, what are the real differentiators, to help you choose your project partners, without compromise on quality?

YOU MAY BE aware of the Aalberts brands, Aalberts hydronic flow control, often represented through their Flamco products, and Aalberts integrated piping systems, often represented through their Pegler and VSH products. The combined Aalberts package gives you a unique opportunity for partnership, allowing you to take advantage of functional savings, like time won on projects, reduced labour costs and overall efficiency improvements, throughout the whole project, from conception to completion.

The Aalberts advantage

The combined Aalberts offer is built around three pillars, helping those delivering complex projects, at scale:

- **The most complete portfolio in the industry**
  The Aalberts portfolio is spread across a data centre building. With plant, valves, pipework and connection technologies, gaps can be eliminated, with one overall supplier, managed by one point of contact – making your life easier

- **End-to-end project execution**
  From early design and specification through prefabrication, logistics and installation support, the Aalberts focus is on taking friction and risk out of the delivery process; fewer suppliers equal fewer coordination risks, plus a faster build versus traditional build
- **Scalable, standardised systems**
  Solutions are engineered to deliver consistent performance across multisite programmes, helping to reduce operational costs, support SG targets and to maximise uptime with repeatable, factorytested designs

From plantroom specialist to fullbuilding partner

Both sides of the Aalberts business have had historical presence within data centres. In and around the plantroom, Aalberts hydronic flow control technologies (such as Flamco pressurisation, expansion and degassing technology) have been present. More recently, prefabricated technical solutions have come to the fore, with the dedicated prefabrication facility in Leszno, Poland becoming a hub for offsite build.

On the Aalberts integrated piping systems side of the business, the focus has been on valves, connection and pipework, throughout the data centre building, connecting elements and keeping critical infrastructure online, with reliable distribution and control of liquids.

The Aalberts portfolio is the widest in the industry. Thanks to this range and technical knowledge, the available scope is broad. With plant and prefabricated elements now sitting alongside valves, connection technologies and piping, Aalberts can support every major fluid system in the building.

Using data centre zones as a basis, this breadth becomes even clearer. From the plantroom, through distribution networks, and into the technology cooling system (TCS) loops serving the data halls, Aalberts supports fluid movement, control and protection at every stage. Each zone brings different demands in terms of pressures, temperature, flow rates and materials, and Aalberts engineer for all of these, incorporating reliability, safety and long-term performance.

Through a holistic approach and delivery of missioncritical hydronic cooling and piping solutions that help to support system reliability, leakfree operation and energy efficiency, across the whole estate, Aalberts is uniquely positioned as "the lifeblood of every modern data centre".

Rather than pulling information, data and documents together from multiple vendors for each subsystem and system, stakeholders can align with a single, expert technical partner, reducing integration effort and procurement complexity.

This breadth of expertise can open up new possibilities. From plantroom to data hall, systems can be connected with minimal intervention, all supported by Aalberts. Rapid connection of without the need for hot works can be enabled. Using the same approach, training can be simplified, as can spares and maintenance.

End-to-end support for a changing market

Data centres are evolving fast as AI workloads, sustainability commitments and regulatory expectations reshape design assumptions. Aalberts recognises that success now depends on more than simply supplying highquality components; it requires end-to-end support, from earlystage concept development to commissioning and optimisation.

Across each zone of the facility, this means selecting solutions that are appropriate not only for today's operating conditions, but for future load profiles and regulatory expectations.

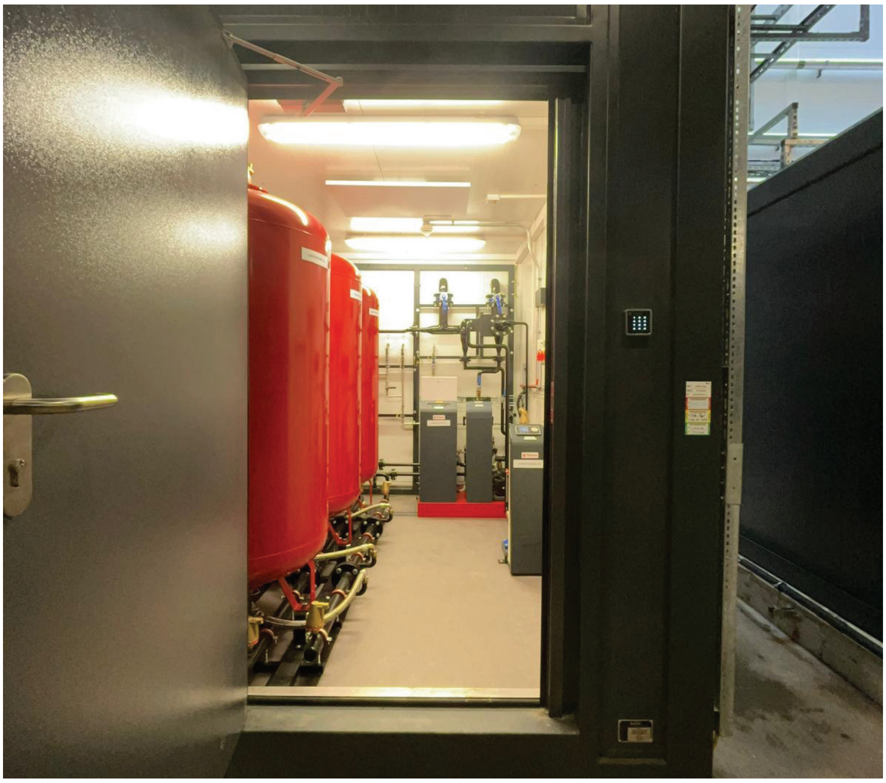

This could manifest in several ways: larger valve sizes in primary plant; different connection strategies in distribution networks, or material and temperature considerations closer to the IT load.

With a comprehensive portfolio that truly spans the whole building, supported by inhouse engineering and factorytested modular solutions, Aalberts is the critical, necessary "lifeblood" behind the next generation of highdensity, lowcarbon and alwayson data centres.

Benefits that matter

Aalberts' data centre proposition focuses on outcomes that directly impact delivery, performance and risk. These benefits fall into two core areas: design and construction and system performance, supported by in-house engineering expertise and customised, factory-tested solutions.

Design and construction

- **Assured compliance and safety**
  Navigating regional regulations across multiple projects can be complex. Aalberts supports project teams early in the planning phase, embedding compliance into the design through proven, factory-tested solutions and established safety standards, which reduce risk and give regulatory confidence to all stakeholders, from Day 1. Aalberts integrated piping systems have recently upgraded and extended their valve range suit both FWS and TCS requirements, reflecting the constant focus on evolution and advancement within the business.

- **Quicker, more predictable build programmes**
  Prefabricated elements enable parallel offsite construction, shortening build times and improving cost certainty compared with traditional, linear installation approaches

- **Scalable, future-ready design**
  Modular platforms make it easier to replicate proven designs, expand capacity and adapt to evolving power densities or cooling technologies. Aalberts' engineering teams support consultants with modelling and system selection to deliver flexible, site-specific solutions without sacrifice

System performance

- **Optimised cooling efficiency**
  Advanced thermal management and precise hydronic balancing reduce wasted pumping energy and maintain stable operating temperatures, improving equipment reliability and extending asset life

- **Lower lifecycle energy and operating costs**
  High-quality materials and optimised system design minimise leakage, pressure losses and thermal inefficiencies, ultimately reducing energy consumption, maintenance needs and total cost of ownership

- **Reliable operation and high uptime**
  By engineering critical systems as an integrated whole, Aalberts reduces single points of failure and maintains service continuity during maintenance or abnormal operating conditions

The element of prefabrication

Prefabricated hydronics solutions are quickly rapidly becoming a strategic differentiator for data centre development teams, under increasing pressure to deliver more capacity, quicker than ever and all the while with less risk. Stakeholders can now work with one partner for resilient cooling and piping solutions across the entire facility, starting with prefabricated plantrooms.

Using a zone-based approach extends naturally into the prefabricated side of things, with assemblies designed specifically for their role in the plantroom, network distribution or TCS environments. This attention to detail means that liquid flow, connection integrity and material selection are optimised, even before equipment ever arrives on site.

- **Speed without compromise**
  Prefabrication within data centre builds has become a practical lever for programme certainty. The integration of pressure maintenance, expansion management and vacuum degassing into factory-built modules that arrive on site pretested and ready for connection to the wider cooling system is a straightforward method for practical savings (time and labour) on site.

This approach shortens onsite MEP activity, simplifies commissioning and reduces the number of trades competing for access in already complex builds.

As fabrication and testing are carried out in controlled environments, quality becomes repeatable rather than dependent on the on-site conditions or labour availability. For operators and consultants, that translates into more predictable PUE outcomes, better protection of water quality and equipment life, and reduced risk of project derailment caused by coordination clashes.

- **Example: Building resilience offsite**
  A recent London colocation project illustrates how prefabricated hydraulic solutions can derisk delivery while meeting stringent local requirements. The new-build data centre, designed around high-capacity chilled water cooling, uses each cooling loop to serve up

Using a zone-based approach extends naturally into the prefabricated side of things, with assemblies designed specifically for their role in the plantroom, network distribution or TCS environments

to 10 MW of IT load, with an N+N configuration. These rooms are fully interconnected, allowing cooperative operation which maximises resilience and operational flexibility, while complying with British electrical and safety standards.

Multiple project challenges demanded a different approach: a compressed delivery programme, constrained plant space and the need for tight coordination between various engineering disciplines.

By resolving interfaces at design stage and offering a prefabricated solution, Aalberts were able to reduce installation and commissioning risk, as well as minimising onsite congestion, which helped the operator maintain build quality and delivery, despite significant time pressure.

Visualising the solution: an immersive 3D experience

To help stakeholders understand how these solutions come together and can be packaged, Aalberts has developed an interactive 3D tour, for virtual exploration.

Accelerate decision making by being able to visually explain the simple yet effective nature of an Aalberts solution. Everything you can see is provided by Aalberts, with 85% of components being manufactured in-house, by Aalberts.

Scan the QR code or follow the link to try the 3D tour for yourself, and experience the Aalberts offer in action!

https://tinyurl.com/aalberts3Dtour


AFL

Pump tech that can keep up with 75–200 kW racks

Why data centers are moving beyond air

THE DAYS when air cooling could handle most data center heat loads are ending. As rack power densities climb into the 75–120 kW range and higher, operators are shifting to hybrid cooling architectures where air becomes secondary and liquid cooling carries most of the thermal load. Direct-to-chip (D2C) liquid cooling is gaining the most momentum because it removes heat at the source with predictable performance and good scalability, even as alternatives like rear-door heat exchangers (RDHx) and immersion cooling continue to grow in specific use cases.

As cooling distribution units (CDUs) evolve to support high-density racks, one component is becoming a defining factor for reliability and efficiency: the pump inside the coolant distribution unit (CDU). Choosing the right pump type isn't a minor design decision, it directly impacts efficiency, responsiveness to load changes, and the economic scalability of the cooling system.

Key takeaways for high-density cooling design

As rack densities rise, the thermal margin shrinks and the cooling system must respond faster and more efficiently. Several points stand out for system designers and operators:

- Liquid cooling is moving from "nice to have" to necessary as densities exceed what air systems can reliably manage.
- Variable load behavior matters. AI and HPC workloads fluctuate, and CDUs must adjust flow efficiently instead of operating at a constant, wasteful point.
- System integration is critical. Pump selection affects cavitation risk, vibration, efficiency losses, and sensor/control strategy.
- The 50–120 kW CDU range is a practical "sweet spot" where compact, high-performance pump design can deliver strong cost/performance balance.
- In-house engineering (design, prototyping, validation) enables precise control over pump curves, materials, and performance envelopes to meet data-center reliability requirements

Why liquid cooling wins

Air cooling has improved steadily, but it is reaching its practical limits at today's heat fluxes. Traditional room air cooling struggles beyond ~15–20 kW per rack; rear-door heat exchangers extend this to ~50–80 kW, but AI/HPC racks at 75–120 kW and beyond generally require direct-to-chip (DLC) liquid cooling.

The advantage comes down to heat transport capability: liquid coolants can carry dramatically more heat than air for the same temperature rise, enabling higher-density computing without the massive airflow, noise, and energy penalties that come with pushing air systems harder. By moving heat removal closer to the chip, D2C systems reduce temperature gradients and help maintain stable inlet temperatures even as rack power rises.

The "sweet spot": designing CDUs for 50–120 kW racks

For CDUs serving modern high-density racks, the pump largely determines whether the loop operates efficiently and safely across real-life conditions. Too much pressure can waste power and accelerate wear, too little flow risks hotspots and instability. In the 50–120 kW range, a well-matched pump can deliver stable operation through rapid load changes while maintaining the pressure/flow characteristics needed for consistent coolant temperatures.

A typical profile for this class of CDU includes:

- **Rack density:** 50–120 kW (often ≥75 kW for AI/HPC)
- **Flow range:** roughly 38–303 L/min (10–80 GPM), depending on ΔT, pressure targets, and coolant properties
- **Coolant:** commonly water/glycol mixtures
- **Design priorities:** low NPSH

requirements, high meantime between failure (MTBF), compact footprint
- **Pump type (in many optimized designs):** centrifugal

Density trends are forcing a new cooling mindset

Rack density has risen quickly over the past two decades moving from single-digit kW racks to today's AI clusters that can exceed 100 kW per rack, with industry projections continuing upward. That growth has outpaced incremental improvements in air cooling and is pushing engineers to optimize every part of the liquid loop.

In this environment, CDU performance affects more than temperature.
It touches uptime, overall facility efficiency, and the ability to standardize and scale deployments. Pump technology sits at the center of that equation because it governs the loop's flow stability and how effectively the system can respond to changing thermal loads.

Centrifugal vs. positive displacement: which pump fits CDUs best?

Two broad pump families dominate industrial fluid systems:
Positive displacement (PD) pumps: Such as gear and diaphragm designs – move a fixed volume per cycle. They are often excellent for high-pressure, low-flow applications. The drawback in many data center CDU scenarios is that PD pumps naturally resist the kind of smooth, efficient flow modulation needed for variable compute loads in relatively low-pressure liquid cooling loops.

Centrifugal pumps, by contrast, use an impeller to convert rotational energy into fluid velocity. They tend to adapt more naturally to changes in system resistance, which can be a strong fit for CDUs that must handle fluctuating demand across racks and manifolds.

In the flow ranges common to high-density D2C systems, centrifugal designs are often favored for:
- Smooth high-volume flow, supporting uniform temperature control across many cold plates
- Strong efficiency with low-viscosity fluids, lowering power draw and operating cost

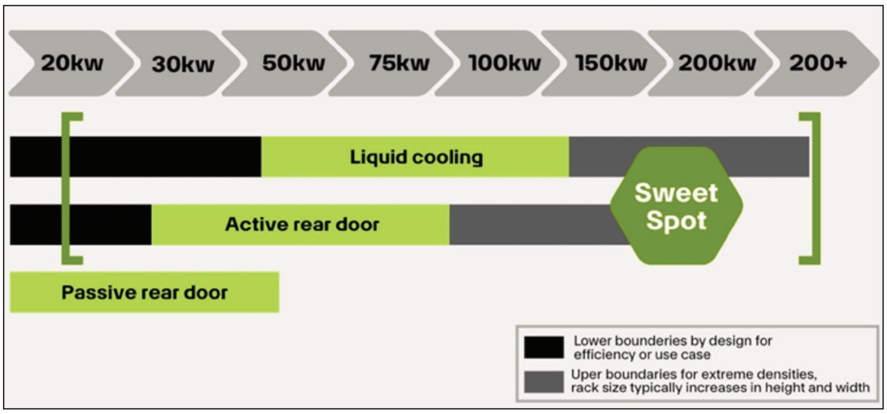

➤ Comparison of Cooling Types

- High performance across typical CDU flow ranges (roughly 10–80 GPMin many configurations)
- Low maintenance and high reliability due to fewer wear-critical moving parts and generally lower vibration

For CDUs, the pump is not just about moving coolant. It's a control element that determines how precisely the system can match cooling delivery to real workload behavior.

Engineering and reliability considerations

As CDU designs scale, reliability becomes a design constraint, not an afterthought. Several engineering factors repeatedly show up as critical in high-density CDU pump selection and integration:
- **Cavitation prevention:** Inlet geometry and operating margins must reduce the chance of vapor formation in low-pressure zones.
- **Material compatibility:** Water/glycol environments demand materials that resist corrosion and wear over long service life (often including stainless steels, polymers, and ceramics).
- **Compact form factor:** CDU footprints are constrained by cabinet space, service access, and the push for higher rack-level density.
- **Thermal endurance:** Pumps must operate continuously at elevated coolant temperatures for years, often under cycling conditions.
- **Scalable output:** Modular approaches that cover a wide flow range help standardize platforms across multiple rack densities.

Tark Thermal Solutions pump modules are qualified for continuous operation under load, pressure, and temperature cycling representative of Tier 3+ data center environments.
Designing for efficiency: operating near the "best point"

One practical efficiency goal in liquid cooling loops is keeping pumps near their best efficiency point (BEP) for the majority of operating hours. When pumps are oversized or forced to run far from their efficient region, energy consumption rises and mechanical wear can increase. A common strategy is right-sizing pumps, so they operate near BEP for most duty hours while maintaining a safety margin for NPSH and transient conditions.

What does this mean for next-generation racks

As data centers move into a future of 75–200 kW racks and beyond, liquid cooling is becoming foundational infrastructure. Within that shift, CDU pump selection and integration will increasingly determine whether systems scale efficiently and reliably.
Centrifugal pump approaches, when correctly engineered for the operating envelope, can offer the adaptive control, efficiency, and service life needed to keep high-density compute stable under rapidly changing loads.

The centrifugal water pumps by Tark Thermal Solutions are engineered to deliver high-performance liquid cooling for next-generation data center enclosures and coolant distribution units. Designed for 50–150 kW data center applications, the latest from the series, TCDU150, helps operators manage rising rack power densities while improving efficiency, reliability, and sustainability.

Learn more at:
tark-solutions.com

Why Essentra Components? Tailored solutions for the data centre industry

At Essentra, we know that even the smallest components can have a major impact on data centre efficiency, security, and longevity. As a global leader in the manufacture and supply of essential components, we provide an extensive range of high-quality standard parts backed by the ability to deliver fully tailored solutions when off-the-shelf simply isn't enough.

FOR the third consecutive year, Essentra will be exhibiting at Data Centre World, where our expert team will be on hand to discuss industry challenges and showcase our latest electronic locking systems. Designed for modern data centre environments. These solutions deliver enhanced security, control, and reliability for critical infrastructure.

In-house expertise for unmatched quality

Essentra's strength lies in our extensive in-house tooling, manufacturing, and testing capabilities. With 14 manufacturing facilities worldwide, we deliver precision-engineered components that meet the highest industry standards. Our tooling teams work closely with customers to optimise designs for performance and durability, while our R&D teams conduct rigorous testing including IP sealing for water and dust, salt spray corrosion testing, and environmental durability assessments, all carried out in-house to IP specifications.

Strategic acquisitions have further strengthened our offering. The acquisition of Hengzhu, a leading Chinese manufacturer of locks, latches, hinges, and handles, enhances our capabilities in sectors such as electric power and telecoms. Meanwhile, Mesan in Turkey expands our global reach and high-quality hardware portfolio ensuring we can support customers wherever they operate.

Broad product range: Quality you can rely on

With a portfolio of 45,000 products and over 1 billion parts in stock, Essentra ensures fast availability from local inventory, helping to minimise downtime.

Our comprehensive range covers everything from IT cabinet hardware to sealing and cable management solutions, including:

- Locking solutions, including swing handles, locks and rotary latch systems
- Hinges for doors and panels
- Feet and castors
- Sealing gaskets and edge protection
- Handles
- Cable entry and cable management
- Fasteners & Fixings
- PCB hardware
- Precision fasteners
- Motion control

Our recently expanded ranges reflects our commitment to innovation, offering flexibility in materials, finishes, and performance requirements, including high-strength stainless steel options for demanding environments.

Custom solutions: Designed around you

Data centres often require more than standard components. That's why Essentra offers tailored solutions designed around your application:

- Custom colours and branding, including powder coating, engraved logos, and printed branding
- Specialised or bespoke components, from minor product modifications to fully custom-engineered designs

Our experts work closely with you to ensure every solution aligns with your operational and security requirements.

Global reach, local support

Operating in 29 countries with 24 strategically located distribution centres, Essentra provides fast, reliable delivery backed by local technical support. Working with over 70% of the world's global manufacturers, we are a trusted partner to leading operators worldwide.

Designed for performance and sustainability

Essentra components are engineered to enhance data centre lifecycle performance, from sealing solutions that prevent dust and water ingress to durable finishes that withstand harsh environments. We also support sustainability goals through the use of recyclable materials and by incorporating up to 50% recycled content across many of our protection ranges.

Why choose Essentra?

With 70+ years of manufacturing expertise, Essentra combines scale and innovation to support data centres around the world. From off-the-shelf products to bespoke engineered solutions, we're here to deliver.

Visit Essentra at Data Centre World Stand D180 to discover our latest innovations, including our new electronic locking systems, or get in touch to discuss your requirements.
T: 0345 528 0474
E: sales@essentracomponents.co.uk

Tudertechnica: high-performance technical hoses for data centre cooling

The global infrastructure of the digital economy is undergoing a fundamental transformation. As artificial intelligence (AI), machine learning, and high-performance computing (HPC) push server densities to unprecedented levels, traditional air-cooling methods have reached their physical and economic limits. In this high-density landscape, cooling systems are essential for maintaining optimal temperatures and ensuring the performance and efficiency of data centre equipment.

BY ANNA PRIMON, APPLICATION ENGINEER, TUDERTECHNICA

TUDERTECHNICA, an internationally recognized leader in the design and manufacture of specialized technical hoses, is shifting new standards for the data centre industry. Leveraging decades of expertise, the Company provides reliable flexible solutions to be applied to Coolant Distribution Units (CDU), chillers, and manifolds within liquid cooling systems. Drawing on Italian manufacturing excellence and on a global distribution network, Tudertechnica delivers hoses with advanced safety standards, ensuring that high-density computing environments can scale sustainably and securely.

The liquid cooling paradigm shift
As AI-driven workloads and High-Performance Computing (HPC) continue to push server densities beyond the limits of air cooling, liquid cooling has become the industry standard. While hybrid approaches combining air and liquid cooling are common, water-based cooling is particularly effective.

In such cases, the combination of air and liquid cooling systems can be a highly effective solution, particularly when dealing with high-heat densities. There are several data centre cooling systems used: Computer Room Air Conditioning (CRAC), Computer Room Air Handler (CRAH), hot and cold aisle containment, in-row and in-rack cooling, chilled water, liquid cooling, free cooling, evaporative cooling, etc.

Liquid cooling involves using a fluid - usually water or a specially formulated coolant - to absorb and remove heat from data centre equipment. Within the industry, these systems generally fall into two primary categories:

- **Direct-to-chip cooling**
In this approach, water-based coolants such as PG25 circulate through tubes that are in direct contact with the processor throughout the cold plate, enabling effective dissipation of heat. The system provides efficient cooling, making it suitable for high-density applications.
- **Immersion cooling**
In this system servers or components are entirely submerged in dielectric (non-conductive) oil. The heat generated by the components is transferred directly to the fluid, which a heat exchanger then cools. Immersion cooling can achieve significant cooling efficiency in high-performance environments.

Benefits of liquid cooling in modern data centres

Liquid cooling offers several advantages over traditional air cooling, making it an attractive option for modern data centres. Here are some of the key benefits:

- **Increased efficiency**
Thanks to liquids, a higher thermal conductivity is attainable. This improved efficiency enables data centres to achieve lower temperatures with less energy, reducing the overall power usage effectiveness (PUE).
- **Reduced energy consumption**
Liquid cooling systems can maintain optimal temperatures with less reliance on mechanical cooling (e.g. air conditioners or fans), reducing the energy consumption associated with cooling infrastructure. This not only lowers operational costs but also contributes to sustainability goals.
- **Space savings**
Since liquid cooling systems can more efficiently manage higher heat densities, data centres can achieve greater computing power within a smaller footprint. This translates to significant space savings, especially in edge data centres or facilities with limited room for expansion.

Liquid cooling efficiency: why the CDU is essential

Liquid cooling is primarily managed by a Coolant Distribution Unit (CDU), which facilitates heat transfer to a chiller, dry cooler, cooling tower, or a combination of these systems. By integrating air and liquid methods, hybrid cooling ensures precise temperature regulation within high-performance data centres.

As the most common implementation for both Direct-to-Chip and Immersion Cooling, the CDU acts as the vital interface between the primary and secondary cooling loops. This separation is essential for preventing condensation within IT Equipment (ITE), maintaining specific coolant chemistry across loops, and providing flexible temperature control tailored to the hardware's needs.

Strategic benefits of CDU integration

By utilizing a liquid-to-liquid heat exchanger, the CDU efficiently rejects heat from the ITE while isolating sensitive components from less regulated Facility Water Systems (FWS) or Chilled Water Systems (CHWS). The Cooling Distribution Unit (CDU) transfers heat from the Technology Cooling System (TCS) to the Facility Water System (FWS) and typically delivers coolant to several IT racks or cabinets.

There are multiple key advantages, including:

- **CONDENSATION CONTROL:** delivering coolant to the rack or electronics strictly above the dew point.
- **RISK MITIGATION:** isolating electronics from harsh facility water and minimizing coolant volume near the hardware to limit the impact of potential leaks.
- **OPERATIONAL FLEXIBILITY:** providing precise control over coolant temperature, pressure, and purity, parameters that a standard CHWS cannot customize for individual racks.

Optimization through flexible hoses

The integration of flexible hoses within the liquid cooling circuit is a high-value strategy for cost mitigation, simplified design, and ease of service. These hoses are critical for connecting manifolds, CDUs, and chillers. However, selecting the right hose requires rigorous design considerations to ensure long-term compatibility and system reliability.

Material innovation: the HPC portfolio for data centres

Since a flexible hose is in contact with the fluids for an extended period, compatibility with the working fluid must be considered when selecting the hose since fluid composition must be preserved in any case.

A great deal of effort has been invested in verifying the compatibility of the selected polymers with the various fluids used in these applications, each test was performed in compliance to OCP's guidelines (Open Compute Project). The fluid must not be altered in any way to cool the systems effectively and for long periods of time, avoiding leaks, contamination, and failures.

- **EPM (EPDM rubber category) hose solutions:** TURAD® series Peroxide-cured rubber is a reliable polymer for the inner lining hose to meet PG25 coolant solutions, since it ensures high reliability over a wide range of operating temperatures and flexibility. Peroxide curing is paramount to avoid triazoles absorption.
- **SILICONE hose solutions:** TUSIL® series This product has been designed to withstand extreme temperature ranges (-60°C to +200°C) and ensures unmatched flexibility.
- **PTFE Solutions:** TUFLUOR® series Tudertechnica's PTFE hoses combine inertia and absolute compatibility both with PG25 and dielectric oil types.

- **HNBR: TURAD® OIL series**

Is a promising polymer to convey dielectric oil-based fluids. The bending radius of a flexible hose can vary greatly and it is typically the smallest radius at which the hose can be bent. It is essential to design the product in order to ensure that the hose doesn't deform. Fittings also must be well selected and applied to the hose since no leakage must occur and corrosion must be avoided.

Safety and fire resistance: UL 94 V-0 compliance

In terms of safety and fire resistance, the cover hoses meet V-0 criteria according to UL 94:2023, providing an extra layer of protection in high-temperature environments and helping to mitigate the risk of flame propagation. The flexible hoses facilitate installation and routing in confined spaces while reducing mechanical stress on connections. They are tested to withstand the bending and vibration cycles that are typical of data centre environments. The covers are resistant to abrasion, ageing, ozone and UV rays, ensuring long-term durability.

Future-proofing: a constantly expanding portfolio

At Tudertechnica, as a member of the Open Compute Project (OCP) - the world's leading community for open hardware and data centre infrastructure, innovation is a continuous process. We are actively investing in R&D to develop next-generation materials. Our product range is in constant expansion, adapting in real-time to the emerging requirements of AI-ready infrastructures and hyperscale requirements.

Your strategic partner in thermal management

Choosing Tudertechnica means securing industrial reliability. As data centres become the foundation of global progress, our hoses provide the security and efficiency required to ensure the continued operation of these systems. From the cutting-edge manufacturing facility to the server rack, the company delivers quality that only an industry leader can provide.

Should you require a technical consultation or wish to explore our latest innovations in data centre cooling, please contact the Tudertechnica expert team: **info@tudertechnica.com**

AEP Global – Engineering excellence delivered with precision

AEP Global is a specialist engineering and technical solutions provider supporting onshore and offshore industries worldwide.

FOUNDED in 2019, the company was established on a deep understanding of the operational, technical, and safety challenges present within complex and highly regulated sectors. Since its inception, AEP Global has built a strong reputation for delivering structured, project-ready engineering services that enhance performance, reduce risk, and create measurable long-term value for its clients.

Operating across data centres, renewables, oil & gas, petrochemical, and nuclear environments, AEP Global combines technical expertise with disciplined delivery methodologies.

The organisation is structured around specialist project teams, technical leadership, and a commitment to predictable outcomes. Its delivery-focused model allows seamless integration into client operations, supporting projects from early planning and design phases through execution, commissioning, and final completion. Quality, safety, and environmental responsibility form the foundation of AEP Global's operations. The company holds internationally recognised certifications including ISO 9001, ISO 14001, ISO 45001, and Cyber Essentials Plus, reflecting a consistent and responsible approach to governance, compliance, and risk management.

Case studies across multiple sites demonstrate a 100% compliance record, reinforcing the organisation's credibility within mission-critical and regulated industries.

Electrical safety and Safe Systems of Work (SSOW) represent a key differentiator within AEP Global's service portfolio. Through advanced isolation management methodologies and digital permit strategies, the company provides real-time visibility and control across plant and infrastructure environments. This flexible approach allows integration with both traditional paper-based systems and fully digital platforms, ensuring alignment with site-specific safety rules while maintaining operational efficiency.

Testing and commissioning services further strengthen the organisation's technical capability. AEP Global delivers comprehensive commissioning management, documentation creation, integrated system testing, and structured start-up execution across a broad range of sectors.

From factory and site acceptance testing through to field calibration, load banking, and cause-and-effect verification, the company ensures every project is commissioned safely, efficiently, and to the highest engineering standards.

GlobalEye SSOW consolidates permits, isolations, and safety controls into a single digital environment, delivering real-time visibility, accountability, and audit transparency across projects and facilities. The platform enables organisations to transition from paper-based processes to a smarter, more reliable safety framework that enhances compliance, reduces administrative overhead, and strengthens operational control

High-voltage electrical services provide another cornerstone of expertise. With competency spanning 132kV through to low-voltage infrastructure, AEP Global delivers HV cable testing, switching plan development, electrical coordination, and infrastructure optimisation.

These services are designed to minimise downtime, maintain productivity, and ensure full compliance with safety regulations and operational procedures, giving clients confidence in the reliability and resilience of their electrical systems.

The organisation also delivers specialist hazardous-area inspection and compliance services supported by CompEx-certified inspectors and adherence to ATEX, DSEAR, and IEC standards. Tailored inspection programmes, preventative maintenance strategies, and rapid rectification solutions enable clients to maintain safe, compliant installations while reducing operational disruption and long-term risk exposure.

Beyond technical delivery, AEP Global provides consultancy services focused on commissioning management, SSOW permit office implementation, and front-end project support. Its consultants bring extensive cross-sector knowledge combined with a proactive, detail-driven approach that supports clients in meeting critical milestones and delivery deadlines.

Whether through remote advisory services or on-site leadership, the company is recognised for delivering practical, results-oriented solutions aligned with project objectives.

At its core, AEP Global operates as a long-term project partner rather than a short-term resource provider. Each engagement is tailored to the client's operational requirements, with emphasis placed on governance, transparency, and sustainable outcomes.

By aligning technical capability with structured execution and measurable performance, AEP Global continues to position itself as a trusted engineering partner across some of the world's most demanding industries.

GlobalEye SSOW – Intelligent digital safety management

Complementing its engineering and compliance capabilities, AEP Global supports the deployment of GlobalEye SSOW, a fully integrated web-based electronic Permit to Work (ePTW) and Isolation management platform.

Designed by experienced professionals from both onshore and offshore energy sectors, the system has been built from the ground up to reflect real-world operational needs while incorporating recognised industry best practice.

GlobalEye SSOW consolidates permits, isolations, and safety controls into a single digital environment, delivering real-time visibility, accountability, and audit transparency across projects and facilities. The platform enables organisations to transition from paper-based processes to a smarter, more reliable safety framework that enhances compliance, reduces administrative overhead, and strengthens operational control.

By combining practical field expertise with modern technology, GlobalEye SSOW provides a forward-thinking approach to safety and permit management. The result is a scalable, intelligent system that supports safer operations, improved governance, and increased efficiency across complex industrial environments. Together, AEP Global and GlobalEye SSOW represent a unified commitment to engineering precision, digital innovation, and measurable safety performance on a global scale.

2026's Perfect Storm: AI growth vs. power availability

TES Power's Managing Director, Michael Beagan, explores the **top five power challenges** the data centre industry may face in 2026, in light of the surge in AI. With a huge increase in demand, power is becoming the most precious commodity across the data centre industry, and this piece delves into the details of why.

WORLDWIDE DEMAND for digital infrastructure continues to surge, predominantly driven by the AI boom. However, data centre builders looking to capitalise on this generational hunger for greater compute capacity will enter 2026 facing myriad challenges, new and old. AI data centres are consuming extraordinary amounts of power, and AI workloads are starting to reshape the digital infrastructure that hosts them. Increasingly, these trends are coming into conflict with constrained power grids, supply chain woes, and stricter ESG requirements. The upshot is that data centres face unprecedented pressure on their power infrastructure.

In 2026, the industry will have to navigate this convergence of grid limitations, supply chain constraints, and evolving regulatory environments. Highlighting this rapidly evolving and uncertain landscape, we're going to take a look at the five biggest power-related challenges shaping the data centre landscape, and the emerging strategies organisations are using to stay ahead.

T1 Market power constraints

Goldman Sachs Research forecasts global power demand from data centers will increase 50% by 2027 and by as much as 165% by the end of the decade, reflecting the massive worldwide surge in compute capacity required to support the booming AI sector. The problem is that the AI, data centre, and energy industry all move at different speeds.

Data centres can typically be built in one to two years. That's still slower than AI companies would like, but the real challenge is that the electricity infrastructure needed to support new data centres takes much longer to get connected. The timelines required to expand transmission networks, substations, and generation capacity stretch far further into the future. In the EU, wait times for securing a grid connection can range from two to ten years. More and more, this mismatch is becoming a critical bottleneck.

Meanwhile, hyperscale and AI-focused developers continue to cluster near major cities where demand is highest, placing additional pressure on already stressed local grids.

Racks supporting AI training workloads can exceed 40 kW today, trending toward 85 kW, with 200–250 kW per rack projected by 2030, and long-term visions discussed by Google exceeding 1 MW racks. These power densities are running up against hard physical and regulatory limits. Several major hubs, including Dublin and Amsterdam, have already paused new grid connections due to capacity constraints before the AI boom. Increasingly saturated markets in the FLAP-D region are likely to follow suit.

Put simply, Tier 1 markets can no longer support growth at the scale or speed required by AI infrastructure. Developers can't continue building in the same locations in the same ways as before.

In 2026 and beyond, data centre companies need to look seriously at new geographic hubs with direct access to available power. Regions such as the Northern Ireland coast could offer proximity to generation assets and more favourable grid availability. This macro-level strategy helps decongest Tier 1 markets and unlock faster development cycles.

AI training is redefining data-centre power consumption

It's no secret that AI increases the overall energy demand of data centres. However, less well known is the fact it also reshapes how and where power is consumed. AI training workloads require enormous, sustained electricity use but are (a bit) less latency-reliant than cloud or AI inference workloads, enabling facilities to be placed far from population centres. As a result, there's a new generation of data centres dedicated especially to training large AI models that, while physically larger and more power-intensive, can be located in more remote areas. Outside congested, largely urban, hubs, data centres can capture advantages like lower land costs, access to water, and abundant renewable power.

However, building any data centre (let alone one big enough to support mass AI training and similar applications) anywhere off the beaten track presents logistical issues. Doing so creates complexities around construction, manufacturing logistics, and commissioning, and traditional approaches to building in Tier 1 markets often fall short.

Data centres must find ways to deliver massive compute capacity in remote locations while overcoming logistical challenges in construction and deployment. AI-focused sites require a different design and build methodology. Companies that bring modular construction expertise and experience delivering remote, large-scale facilities can help accelerate deployment, reduce risk, and simplify the thorny prospect of building far outside established industrial hubs.

Land scarcity demands greater flexibility

Tied into the first two challenges is the problem of land. Data centres are big business, but more prosaically, they're very big buildings, often with attendant outbuildings for hosting UPS and backup generators. Securing suitable plots for data centre development is becoming increasingly difficult and expensive, especially in or near major metro areas. Sites with the right combination of access to power, fibre, zoning, and environmental compliance are increasingly scarce, driving up costs across multiple markets.

Developers are finding themselves faced with interminable permitting cycles, not to mention growing community scrutiny around noise, water usage, and sustainability.

These land constraints often force operators to make trade-offs: building more capacity in multiple locations with smaller footprints, compromising on ideal site layouts, or locating closer to power generation at the expense of also being close to an increasingly hostile general public.

Flexibility is becoming a key differentiator. Leveraging modular, highly customisable power systems makes it easier to adapt infrastructure to unconventional or constrained sites. This helps reduce the burden of site selection and enables operators to maximise usable capacity within limited or irregular parcels of land.

Tightening cross-border regulations and supply chain pressures

The global data-centre supply chain is already under significant strain due to surge in demand. Now, data centre companies (often international firms with sprawling supply chains) are having to contend with heightening cross-border complexity. Continued post-Brexit fallout is snarling up the border between the UK and Europe, and US tariffs — bad enough on their own — have ushered in a new era of frosty trade relations and more isolationist policies. Setting aside the increased cost of moving construction materials and data centre components around the world, customs complexities are also driving up lead times. Critical components such as switchgear, transformers, batteries, and high-density compute hardware all face growing bottlenecks. Deloitte reported recently that large original equipment manufacturers are increasingly locking themselves into multi-year agreements to produce key components, like transformers, switchgear, power management equipment, and power generation systems with single clients. As a result, these critical components are sold out, with restocks not expected to arrive for several years.

All of these challenges mean steep increases in lead times across core electrical infrastructure. These delays directly threaten deployment schedules for new builds and expansions.

To navigate this environment, data centres are leaning on partners with localised manufacturing, diversified supply networks, and the ability to operate across regulatory borders. Companies with operations in Northern Ireland, for example, enable clients to take advantage of UK-based trade agreements (including with the US) as well as Northern Ireland's unique relationship with the EU offering operators a way to mitigate cross border complexity Partnering with companies that have in-house engineering and assembly capabilities also reduces reliance on constrained international supply lines.

Evolving ESG regulations increase operational and design pressures

Governments are accelerating ESG mandates as net-zero deadlines approach, with the EU typically leading the way. For data centres, this means tighter scrutiny on power efficiency, emissions, water usage, generator runtime, and backup-power resilience.

At the same time, operators must secure access to a growing but still limited pool of clean energy sources. Variability in renewable generation introduces added complexity for UPS systems and backup infrastructure, creating concerns about reliability under fluctuating loads, especially when regulatory frameworks restrict generator size and design specifications.

Adopting modular, easily serviceable power systems allow for rapid replacement and minimal downtime. Predictive monitoring capabilities enhance visibility across power-distribution components, improving both "time to detect" and "time to resolve."

Flexibility and resilience in the face of a changing power landscape

In 2026, the power landscape for data centres will be defined by scarcity, complexity, and accelerating demand from AI. Operators that explore new locations, embrace modularity, rethink supply chains, and keep a watchful eye on evolving ESG requirements will be best positioned to deliver reliable capacity at scale.

Partnering with organisations that have expertise related to remote builds, modular power systems, and cross-border operations is becoming an important competitive edge for data centre companies looking to navigate the challenges to come.

The 2AM Troubleshooting Call

That Could Have Been Prevented

Outdated cable documentation turns routine maintenance into all-hands emergencies because nobody knows what is plugged in where. ***Rapid*ID™ eliminates the guessing game:** pre-labelled patch cords with barcode scanning deliver real-time accuracy.

RAPIDID™

Reduce Time and Cost

of cable documentation compared to traditional manual approaches

Automate Labour-Intensive and Error-Prone Cable Documentation Process

that can lead to an outage

Free Up Network Engineers

to focus on more rewarding and strategic initiatives

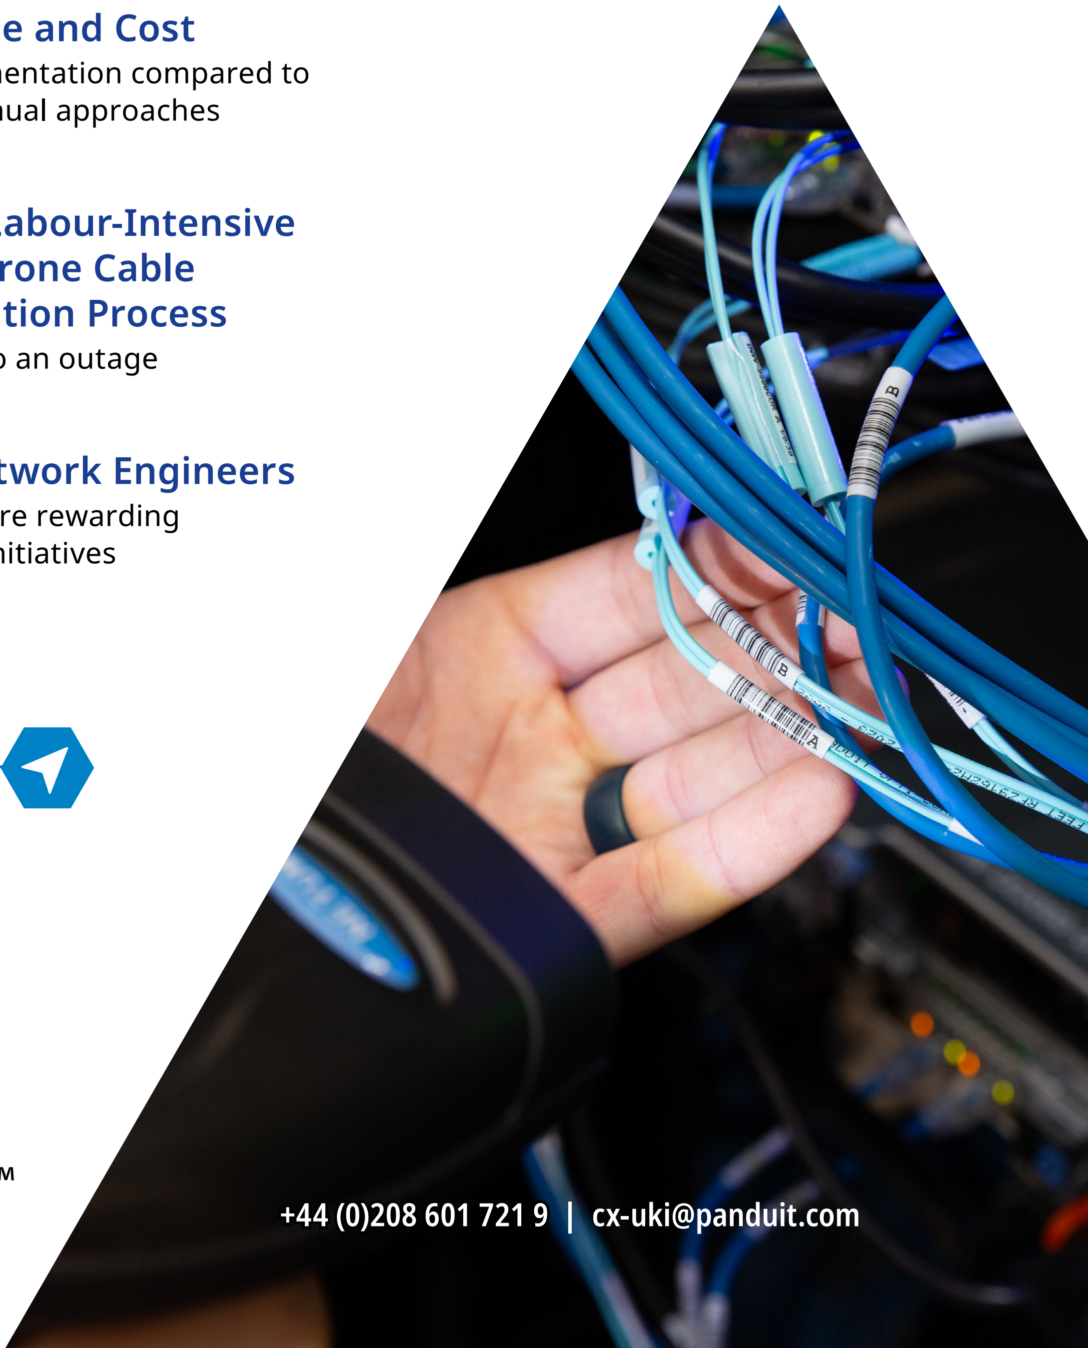

For more information, visit
www.panduit/rapidid

PANDUIT™

+44 (0)208 601 721 9 | cx-uki@panduit.com

Innovation without compromise: The case for digital twins

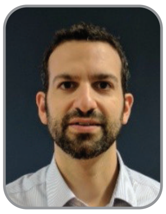

Ultimately, in an industry undergoing scrutiny and accelerating demand, digital twins offer a smarter, more adaptive way forward.

BY MARK FENTON, PRODUCT ENGINEERING DIRECTOR AT CADENCE

POWER GRIDS keep countries ticking, generating and transporting power. However, a staggering 70% of data center facility leaders around the globe say their national power grid is being stretched to its limits. One key driver of this growing pressure is artificial intelligence (AI). As AI adoption accelerates, powering everything from drug discovery to fraud detection and autonomous assistants, the strain on the data centers that support it is becoming harder to ignore. In fact, three-quarters (74%) of data center operators report that their data center is experiencing increased demand from AI-driven workloads.

As AI reshapes industries and intensifies infrastructure demands, data centers must find ways to scale responsibly. That means reducing the load on shared energy and water resources to help safeguard national grids and the communities that depend on them.

The good news is that AI can equally empower data center leaders, too. With tools like digital twins—virtual models of physical operations—operators can better test, monitor, and fine-tune data center performance. Even better, AI-enhanced digital twins offer more predictive insights and smarter automation, helping to reduce AI's environmental impact and boost overall efficiency. It's ultimately a more intelligent approach that brings greater control to an increasingly complex industry.

Uncovering pain points

Before making any changes, data center operators must first determine exactly where efficiency challenges reside. Commonly, this starts with identifying stranded capacity—installed capacity in the data center that cannot be used. Imagine it's like a game of Tetris, where data centers are playing five levels at once, trying to fit all the blocks (systems) into the data center. Operators are constantly trying to slot systems into place, often without full visibility of available capacity. As a result, the facility risks falling short of its intended performance, while inefficiencies quietly rack up costly consequences.

Currently, stranded capacity is reported as a constraint by 29% of data center operators. In addition, 60% of data center leaders say overprovisioning—allocating more resources than necessary—still occurs within their data centers. This is understandable, as they're seeking reliability, especially during peak periods. However, it wastes energy, driving up the data center's footprint and operational costs.

Moreover, high-density servers, which are designed to deliver exceptional performance, can become part of the overprovisioning issue if not managed correctly. Therefore, it's crucial to ensure these systems are running efficiently, minimizing stranded capacity and avoiding over-provisioning wherever possible.

Hailed as a game-changer for data center operations, digital twins play a critical role in tackling stranded capacity and overprovisioning, ensuring resources are allocated correctly. By providing a real-time, virtual model of the data center's infrastructure and operations, they allow operators to visualize power and space utilization with pinpoint accuracy, identifying underused systems and imbalances in workload distribution. When enhanced with AI, digital twins can go even further, using predictive insights to test how to scale operations without disruption to surrounding resources.

Keeping cool

Whilst stranded capacity and overprovisioning contribute to widespread inefficiencies, cooling is another area of potential concern. Here, data center operators should rethink legacy practices. This is where digital twins continue to come into their own. By creating simulations of the physical infrastructure with digital twins and AI, operators gain a powerful tool when planning for new and improved cooling systems. With the immense amount of energy and heat generated by AI-intensive operations, traditional air cooling is being pushed to its limits. Operators are now turning to liquid cooling, which can manage higher thermal loads more efficiently.

Currently, 45% of data center decision-makers have adopted liquid cooling, with another 19% planning to introduce it within the next year. This highlights a shift driven by increasingly dense server racks and AI workloads that exceed the thresholds of conventional systems. While air cooling handles loads up to around 20kW per rack, hybrid approaches that incorporate liquid cooling are far better suited for racks exceeding 25kW.

Whilst stranded capacity and overprovisioning contribute to widespread inefficiencies, cooling is another area of potential concern. Here, data center operators should rethink legacy practices

Similarly to visualizing power, digital twins enable operators to simulate these systems in advance. As a result, they offer visibility into variables like cooling performance, airflow, and heat recovery potential that would otherwise be a challenge. This helps identify the most effective approach for each facility's specific environment, reducing guesswork and risk. What's more, operators can test multiple cooling scenarios, compare performance, and make informed investment decisions all before physical implementation.

Beyond this, digital twins support innovations like waste heat recovery, further contributing to sustainable operations. In this way, they're not just tools for optimization, they're enablers of more responsible innovation.

The catalyst for change

Data centers are making clear strides toward operational improvement, but turning ambition into action is often difficult, as it involves highlighting existing areas of inefficiency. AI-powered digital twins are emerging as the industry's breakthrough solution here, and 73% of operators already agree. A good starting point is to use digital twins to identify and address stranded capacity and overprovisioning, optimize existing air cooling systems, and test the impact of implementing liquid cooling. By unlocking greater visibility and control, operators can modernize and maximize the efficiency of existing infrastructure before investing in additional capacity.

Ultimately, in an industry undergoing scrutiny and accelerating demand, digital twins offer a smarter, more adaptive way forward. As data centers navigate the realities of the AI era, those that embrace this kind of intelligent, responsible innovation will be best positioned to lead the way.

AEP GLOBAL

GLOBAL EYE

Supplying the Best Globally

AEP Global is a specialist provider of Safe Systems of Work (SSOW) across onshore and offshore industries, supporting projects from construction and commissioning through to operations and long-term maintenance sustainability.

AEP Global provides turnkey solutions including:

- **Fully managed Permit Office teams**
- **Electronic Permit to Work (ePTW) systems**
- **Isolation management and LOTO control**
- **Integrated safety governance processes**

Our proprietary platform, GlobalEye SSOW, underpins our delivery model providing digital Permit to Work and isolation software designed specifically for high risk, high value operational environments.

We combine practical site experience with robust digital systems to ensure safety, compliance, and operational continuity across complex, fast-paced projects.

www. aep-global.com

www. globaleye-ssow.com

The future of AI depends on data centre efficiency

Global data centre electricity demand is set to double by 2030 . By way of comparison, this is the equivalent to adding the entire energy consumption of the UK today to the market in the space of five years. As we look ahead, the curve is only getting steeper. And AI is the reason.

BY SADIQ SYED, SVP DIGITAL BUILDINGS, SCHNEIDER ELECTRIC

AS THE RACE to build bigger, smarter models heats up, a new scramble is underway for the power to fuel them. But there is a scarcity issue that is a very real and present threat. Bridging the gap between the demands of AI and powering the data centres that enable it is critical to its future.

If the global economy is to realise the benefits from energy-intensive technologies like AI, then data centres need sufficient power to operate. This is why countries worldwide are exploring options like increasing infrastructure investment into solar, nuclear and wind power. But these investments can take years to show impact – wait times for securing a grid connection in the EU range from two to 10 years.

Alternatively, operators are building more sites, incurring huge investment in an attempt to get ahead of the power shortage.

To address the energy requirements of AI, we first have to look at our existing data centres. Buildings waste nearly 40% of the energy they use, so being more efficient with what we already have could be a faster fix. As data centres are also under pressure from regulators, local communities and investors to operate more efficiently, it's a win-win for operators.

To do this, they need to understand how poor energy management affects them, what blockers stand in the way of their transformation and how reducing wastage can support businesses looking to get the most out of AI.

An inflection point for business efficiency

Poor energy management is a silent killer. It doesn't just affect the environment but erodes a company's resources. This problem is even more acute within data centres. Facilities often have multiple power supply systems to ensure uninterruptible service, cooling systems, temperature sensors, lighting, both physical and digital security, just to function.

All too often these systems are siloed, making it difficult to get a realistic picture of how the data centre is functioning. Without a unified view of all the systems, the chances increase that engineers miss a voltage imbalance that damages equipment.

Overly complex and fragmented systems can also expose organisations to higher prices. Many utility providers calculate bills based on energy charges, the total electricity used over a month, and demand charges, based on the highest rate of power consumed during any short interval.

If data centre operators don't have full visibility over their systems, they could miss the opportunity to use cheaper solar energy instead of grid power or use several energy intensive systems during peak times due to poor coordination between building and electrical teams.

Given a 100 kw data centre can face over £200,000 per year in electricity costs minimising times when this happens could save thousands of pounds over the course of a year.

Helping engineers help you

Facilities engineers need a high level of technical knowledge to carry out their work. Understanding the information coming from power and energy systems and knowing what to change to reduce waste can take decades to learn properly. However, we have an aging workforce without enough skilled engineers to replace them. And systems are getting more complex, requiring a concrete understanding of the data flowing through the data centre ecosystem.

AI is playing a growing role in enabling data centre engineers. By applying its analytical powers to a platform pooling the disparate systems, it can help engineers translate information into efficiency and empower the next generation with insights to meet stringent compliance targets.

Similarly, by analysing previous data patterns it can help predict issues before they become an actual problem, enabling facility managers to become more proactive and reduce unnecessary damage or downtime to equipment.

At a time when AI is dominating the consumer market too, this generation expects automation to support them at work. To enable the workforce, data centres can be no different.

The foundation of growth is efficiency

Simplicity is the key when it comes to reducing energy wastage. By bringing together the information generated by energy-draining electrical and mechanical systems, with insights on power flowing from the grid, operators can anticipate failures, prevent downtime and extract more performance from the same footprint.

More data centres might well be necessary to meet AI-centric goals of tomorrow, but we must find a way to get us there by future-readying existing sites and the underlying infrastructure.

That primarily means democratising access to disparate systems, so that they don't run in isolation and stay ahead of issues.

If data centre operators want to avoid unnecessarily high energy costs – a blot on the copybook of AI's potential – and the challenges of constant site expansion, optimising existing infrastructure can go a long way.

There is only finite space, and so the next race for data centre operators will be towards simplicity and efficiency, whether by simplifying infrastructure or supporting engineers.

Powering data centres with precision, performance, and reliability

In today's data-driven world, every second of downtime costs money, reputation, and opportunity. At Flowtech, we understand that a data centre is only as strong as the systems supporting it. That's why we deliver innovative engineering solutions, robust products, and expert support that keep critical infrastructure running smoothly and efficiently.

FROM COOLING SYSTEMS and pipework to valves and leak detection technology, Flowtech equips data centres with the tools and services they need to maintain uptime, reduce operational risk, and optimise performance.

Expertise at scale

Flowtech is more than a supplier—we're a partner in operational excellence. With over 620 highly skilled technical specialists across 21 operational locations in the UK, Ireland, and Benelux, we combine global experience with local knowledge. Our 121 engineering experts are trained in fluid conveyance, process automation, and thermal management, ensuring solutions are designed, installed, and maintained with precision.

Our capabilities extend from custom engineering solutions and repair and overhaul services to high-volume bid capacity. With 75,000 products in stock and a 97% availability rate, we ensure our clients get the parts they need – when they need them.

Streamlined product solutions

Flowtech offers an extensive portfolio of products tailored to the unique demands of data centres. Whether it's high-performance valves, flexible pipe systems, or advanced leak detection, each solution is designed to enhance reliability and efficiency.

- **Pipe Systems:** Our BRUGG pipe systems, including Flexwell Safety Pipe for fuel supply and Secon-X for liquid cooling, combine flexibility with monitoring capabilities. These systems reduce installation time, limit downtime, and allow continuous leak monitoring while in operation.
- **Valves:** We provide a comprehensive range of process valves – from floating and trunnion mounted ball valves to high-performance butterfly, gate, globe, and check valves. Each product is engineered for longevity, precision, and zero leakage performance, ensuring data centre cooling and fire suppression systems operate flawlessly.
- **Leak Detection:** SGB leak detectors offer permanent monitoring for double-walled tanks and pipelines. Early detection safeguards critical IT equipment, protects assets, and ensures optimal operational conditions. Our vacuum and pressure leak detectors are robust, reliable, and easily integrated into existing systems.
- **Thermal Management:** For liquid cooling, we supply Parker and Danfoss Hansen solutions, including EPDM-P hoses, quick-connect couplings, and fully integrated cooling systems. These products enhance fluid purity, energy efficiency, and operational safety.
- **Process Automation:** Flowtech's actuated valves, both electric and pneumatic, provide live data, precise flow control, and remote management capabilities. Integrated with BMS and IoT systems, they deliver energy-saving efficiencies, reduce maintenance, and improve overall system performance.

Real-world impact

At Flowtech, we don't just deliver products – we deliver outcomes. Our solutions are tested, trusted, and proven in some of the most demanding data centre environments.

For example, a mechanical contractor in Ireland tasked us with supplying and supporting valves for multiple data centre phases, including backup power systems. Tight schedules and precise specifications demanded expertise. We leveraged our European supplier network to source trunnion-mounted ball valves certified to SIL 3, ATEX, and API 607 standards, ensuring zero leakage performance for critical applications. Regular communication, technical support, and careful scheduling enabled the project to stay on track and meet all deadlines.

In South Wales, we installed 5km of DN100 Secon-X secondary contained pipework across a complex multi-floor data centre. With seven vacuum leak detectors integrated into the system, the project was delivered in just seven weeks. Our flexible, continuous piping installation reduced downtime, allowed for individual pipe monitoring, and gave the client confidence in the reliability of their cooling infrastructure. These case studies illustrate the Flowtech difference: precision, reliability, and a

partner-focused approach that ensures success from specification through installation and ongoing operation.

Enhancing sustainability

Flowtech's commitment extends beyond operational excellence. We actively work to reduce environmental impact through optimised supply chains, efficient fleet operations, and minimal waste strategies. Our smart valve solutions, energy-efficient PICVs (Pressure Independent Control Valves), and advanced monitoring systems support clients in lowering energy consumption and reducing carbon emissions, helping data centres operate sustainably without compromising performance.

Accelerating performance through technology

Innovation is embedded in every Flowtech solution. From IoT-enabled sensors to smart automation, we enable data centres to monitor, control, and optimise their systems in real-time. Our modular PICVs deliver precise flow control while maintaining differential pressure, improving energy efficiency and simplifying commissioning. By integrating Modbus actuators and BMS interfaces, operators can adjust systems remotely, ensuring consistent temperature control and optimal resource usage.

Quick-connect couplings, advanced hoses, and thermal management components are designed for efficiency, safety, and reliability. With flat-face designs that prevent fluid loss, high resistance to vibrations, and reduced pressure drop, these solutions help maintain operational excellence even in high-density, high-power environments.
A Complete, Integrated Approach
Flowtech's strength lies in integration.

We combine engineering expertise, product innovation, and installation proficiency to offer complete solutions. Whether it's designing flexible piping, supplying actuated valve packages, or commissioning pre-functional testing, we deliver end-to-end support that streamlines projects, mitigates risk, and reduces total cost of ownership.

Our collaborative approach ensures that data centre operators, project engineers, and contractors receive actionable solutions that meet the needs of complex and evolving environments. Every project benefits from our direct engagement, technical insights, and proactive problem-solving – so deadlines are met, systems operate reliably, and downtime is minimised.

Trusted by leading global brands

Flowtech's reputation is built on reliability. Our products are installed in Tier 3 and Tier 4 data centres worldwide, supporting cooling, fire suppression, and backup power systems. Leading global brands trust Flowtech to deliver solutions that combine durability, precision, and operational efficiency.

Whether it's supplying high-performance ball valves, resilient seated butterfly valves, or advanced leak detection systems, Flowtech ensures every component meets stringent quality standards and regulatory requirements. Our engineering teams provide guidance, integration support, and customised solutions to meet each client's unique requirements.

Why Flowtech?

Choosing Flowtech means selecting a partner who:

- Delivers reliable, high-performance solutions for critical data centre systems
- Provides in-house engineering expertise and global supplier access
- Reduces downtime, enhances energy efficiency, and lowers operational risk
- Integrates smart automation, IoT, and BMS systems for real-time control
- Offers end-to-end support from design through commissioning and maintenance
- Champions sustainability and reduces carbon footprint across operations

Our approach is simple: we energise, mobilise, and activate solutions that enhance operational performance, protect valuable assets, and enable data centres to thrive.

Looking ahead

The data centre industry continues to evolve at pace. Higher power densities, advanced cooling requirements, and stricter environmental targets demand innovative solutions and expert guidance. Flowtech is ready to meet these challenges head-on.

By combining precision-engineered products, cutting-edge technology, and skilled engineering teams, we enable data centre operators to anticipate problems, prevent downtime, and deliver optimal performance. Our solutions are not just about equipment – they are about outcomes. Reduced energy use, lower maintenance, enhanced reliability, and peace of mind for operators.

Experience Flowtech at DCW
Flowtech will be showcasing our latest data centre solutions at this year's Data Centre World event. Visitors can explore our portfolio, see live demonstrations, and learn how our products and services can enhance uptime, reduce failure, and optimise efficiency across all areas of operation.

From process valves and actuated systems to leak detection and liquid cooling, Flowtech provides a complete, integrated approach that supports operational excellence and drives measurable results. Our team will be available to discuss specific challenges, demonstrate products, and provide insights tailored to your data centre's needs.

Your partner for precision, performance, and reliability

Data centres are the backbone of the modern digital economy. They demand solutions that are robust, reliable, and efficient. Flowtech delivers all three – alongside the expertise, insight, and support to make every system operate at its best.

For operators looking to reduce downtime, improve energy efficiency, and safeguard critical infrastructure, Flowtech provides the products, services, and expertise to make it happen. From valves and piping to automation and leak detection, we offer integrated solutions designed to deliver outcomes that matter.

Visit us at DCW stand C1 to see how Flowtech energises, activates, and transforms data centre operations with precision-engineered solutions and expert guidance. With Flowtech, your infrastructure runs smoother, safer, and smarter.

Scan to see Flowtech's Data Centre projects.


flamco

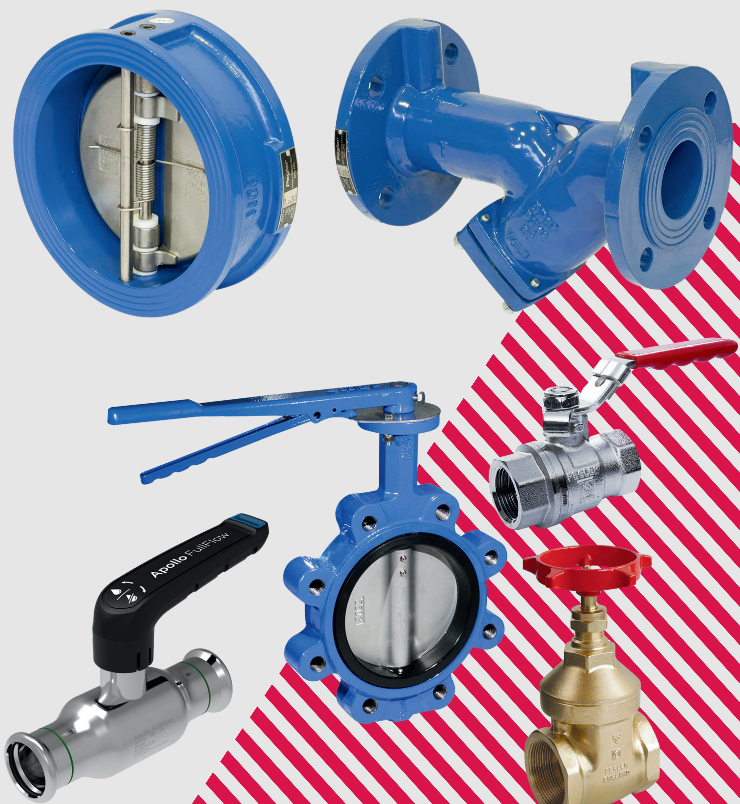
Apollo FullFlow

FibrePath® – Reengineering cable containment for high-density infrastructure

As data centre densities increase and fibre counts continue to climb, cable containment is no longer a passive infrastructure element. It plays a direct role in airflow efficiency, fire safety compliance, operational resilience and long-term scalability.

HIGH-PERFORMANCE IT environments — from hyperscale campuses to edge facilities — are pushing structured cabling systems to new limits. AI workloads, 400G and 800G optical deployments, and increasingly modular architectures demand containment solutions that are robust, adaptable and engineered for precision.

FibrePath®, developed by XSpaceUK, has been designed specifically for this new generation of infrastructure. Moving beyond traditional plastic tray systems, FibrePath introduces a steel-based modular containment platform built to support high-density environments where safety, performance and future expansion are critical.

Rather than functioning as a simple cable pathway, FibrePath forms part of the structural backbone of modern network architecture.

Engineering for high-density environments

Legacy containment solutions were often designed for lower cable volumes and static layouts. Today's facilities are dynamic. Fibre runs increase rapidly, rack configurations change, and interconnectivity requirements expand.

FibrePath addresses these realities through a structured, modular containment architecture. Its systemised design enables clean routing, controlled cable separation and consistent pathway definition — reducing the risk of congestion and improving long-term manageability.

For designers and operators, this provides more than visual organisation. It delivers predictable cable routing behaviour, easier fault tracing and improved upgrade efficiency.

In environments where downtime carries significant cost, clarity and accessibility are operational advantages.

Modular deployment and installation efficiency

Time-to-deployment remains a critical metric in infrastructure projects. Containment systems that are complex to assemble or difficult to modify can introduce delays and unnecessary labour costs.

FibrePath has been engineered for streamlined installation. Modular sections connect securely and intuitively, allowing containment runs to be assembled rapidly while maintaining structural integrity. Standardised fittings reduce on-site fabrication requirements and help maintain consistency across large deployments.

This modularity also supports phased infrastructure build-outs. As additional racks, pods or network zones are introduced, FibrePath can expand without requiring removal or reconfiguration of existing containment routes. For operators managing live environments, this flexibility reduces disruption and preserves operational continuity.

Airflow-conscious design for thermal stability

Thermal performance is directly

influenced by how cabling is managed. Poorly structured containment can obstruct airflow, create hot spots and increase cooling demand.

FibrePath incorporates vented steel channel sections that promote vertical and lateral airflow movement. By reducing cable bundling pressure and enabling more open containment pathways, the system supports cold aisle containment strategies and more consistent temperature control across rack rows.

Improved airflow management contributes to lower cooling energy demand and supports equipment longevity — particularly in high-density compute environments where thermal margins are tight.

As sustainability and operational efficiency become increasingly linked, containment design plays a subtle but meaningful role in overall facility performance.

Fire safety and non-combustible construction

Fire risk mitigation is a non-negotiable consideration in mission-critical infrastructure.

Traditional polymer-based containment systems can contribute to fire spread and generate toxic smoke under extreme conditions. FibrePath addresses this risk through non-combustible steel construction, providing enhanced fire resilience and greater compliance confidence.

Steel containment does not propagate flame and does not emit harmful combustion gases, making it particularly suitable for enclosed data environments where evacuation complexity and equipment density heighten safety requirements.

For operators working within stringent regulatory frameworks, material selection at the containment level becomes a strategic decision rather than a cosmetic one.

Scalability built into the architecture

Infrastructure rarely remains static. Fibre counts increase. Equipment footprints shift. Interconnect models evolve. FibrePath's architecture supports incremental expansion without compromising structural integrity or organisational clarity. Modular sections allow extensions, directional changes and capacity increases to be implemented efficiently.

FibrePath®, developed by XSpaceUK, has been designed specifically for this new generation of infrastructure. Moving beyond traditional plastic tray systems, FibrePath introduces a steel-based modular containment platform built to support high-density environments where safety, performance and future expansion are critical

This scalability is particularly relevant in:

- Hyperscale and colocation facilities
- Enterprise data centres undergoing phased upgrades
- Edge deployments requiring compact yet adaptable routing
- Telecommunications switching environments

By embedding flexibility into the containment layer, FibrePath supports long-term infrastructure evolution without forcing wholesale replacement.

Durability and lifecycle sustainability

Material durability directly influences total cost of ownership. Steel containment provides mechanical strength, resistance to deformation and long service life under continuous operational load.

In addition to structural advantages, steel offers strong sustainability credentials. It is highly recyclable and aligns with circular economy principles increasingly adopted across infrastructure sectors.

Selecting recyclable, long-life containment reduces lifecycle waste and supports carbon-reduction strategies without sacrificing performance or compliance.

In a market where ESG considerations are shaping procurement decisions, responsible material selection carries strategic value.

Operational benefits beyond installation

Cable containment affects day-to-day operations more than is often acknowledged.

Well-defined routing improves fault tracing and speeds maintenance interventions.

Clear pathway segmentation reduces accidental disturbance during upgrades.

Structured containment simplifies documentation and supports compliance audits.

By reducing cable strain, preventing overcrowding and maintaining controlled bend radii, FibrePath contributes to preserving the performance integrity of fibre and copper systems alike.

For infrastructure teams managing complex environments, these operational efficiencies translate into measurable long-term value.

A strategic infrastructure decision

Containment systems are often treated as secondary considerations in early design stages. However, as density increases and compliance requirements tighten, their impact becomes more visible.

FibrePath positions containment as a strategic infrastructure layer — one that supports airflow optimisation, fire resilience, scalability, sustainability and operational clarity simultaneously.

By combining modular engineering with non-combustible construction, FibrePath offers an alternative to legacy containment approaches that were not designed for today's performance demands.

To learn more visit **xspaceuk.com**

The 2AM Troubleshooting Call

That Could Have Been Prevented

Outdated cable documentation turns routine maintenance into all-hands emergencies because nobody knows what is plugged in where. ***Rapid*ID™ eliminates the guessing game:** pre-labelled patch cords with barcode scanning deliver real-time accuracy.

RAPIDID™

Reduce Time and Cost

of cable documentation compared to traditional manual approaches

Automate Labour-Intensive and Error-Prone Cable Documentation Process

that can lead to an outage

Free Up Network Engineers

to focus on more rewarding and strategic initiatives

For more information, visit
www.panduit/rapidid

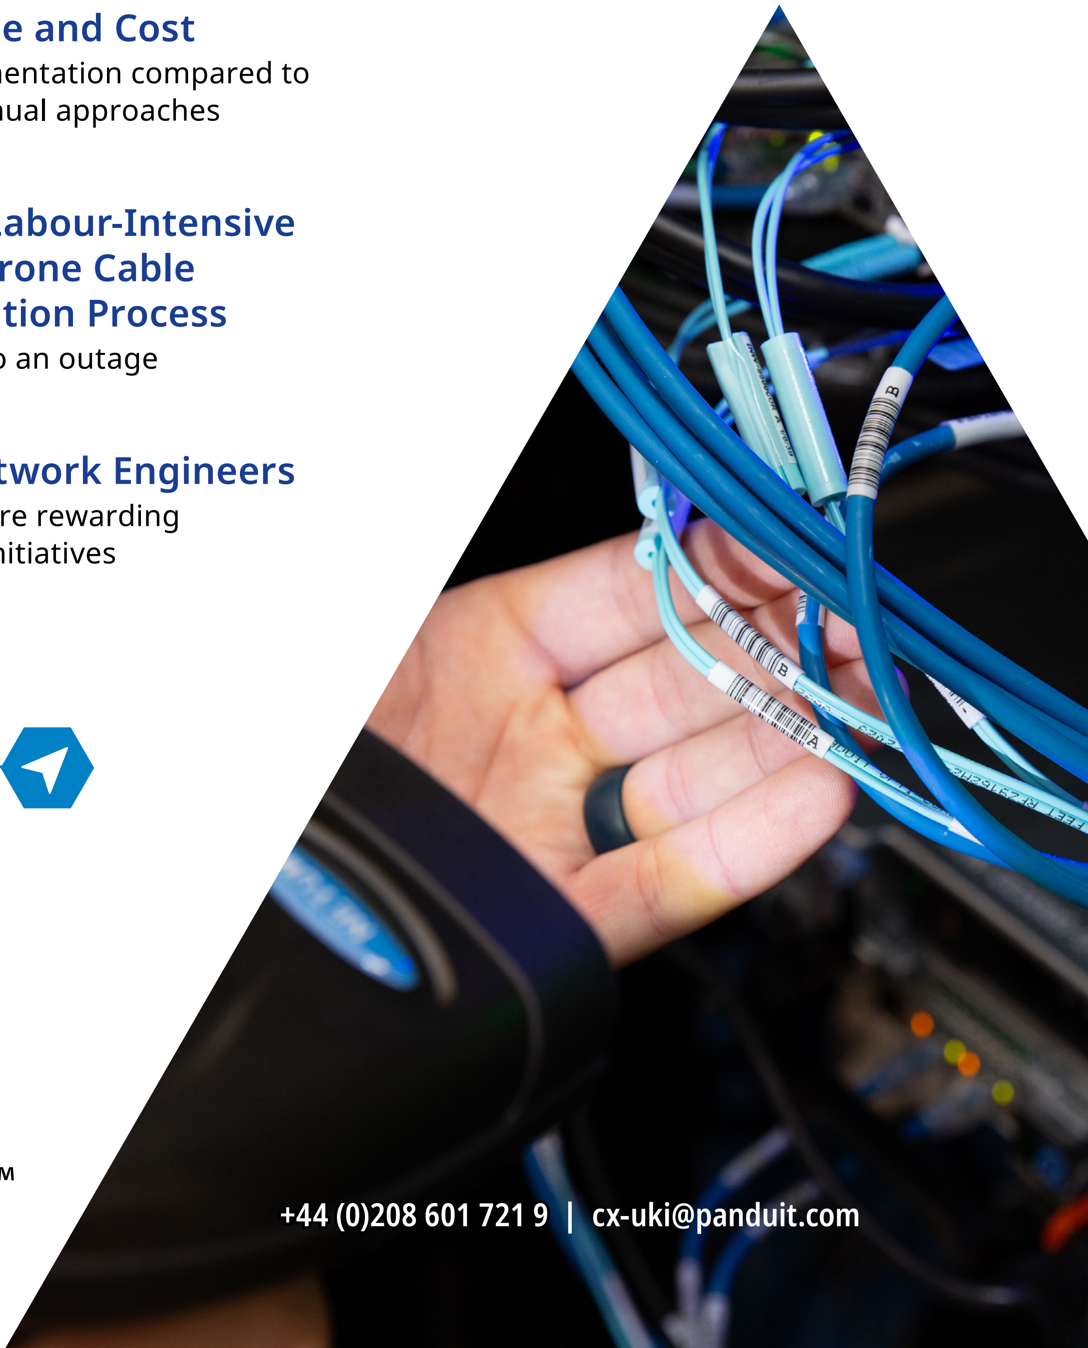

PANDUIT™

+44 (0)208 601 721 9 | cx-uki@panduit.com