# DATACENTRE SOLUTIONS

DCS

SPECIAL EDITION STOCKHOLM 2026

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

HTTPS://DATACENTRE.SOLUTIONS

## WHY EPDS, LCAS AND WEIGHT MATTER IN DATA CENTRE INFRASTRUCTURE

SILVER | Top 15%
ecovadis
Sustainability Rating
JUL 2025

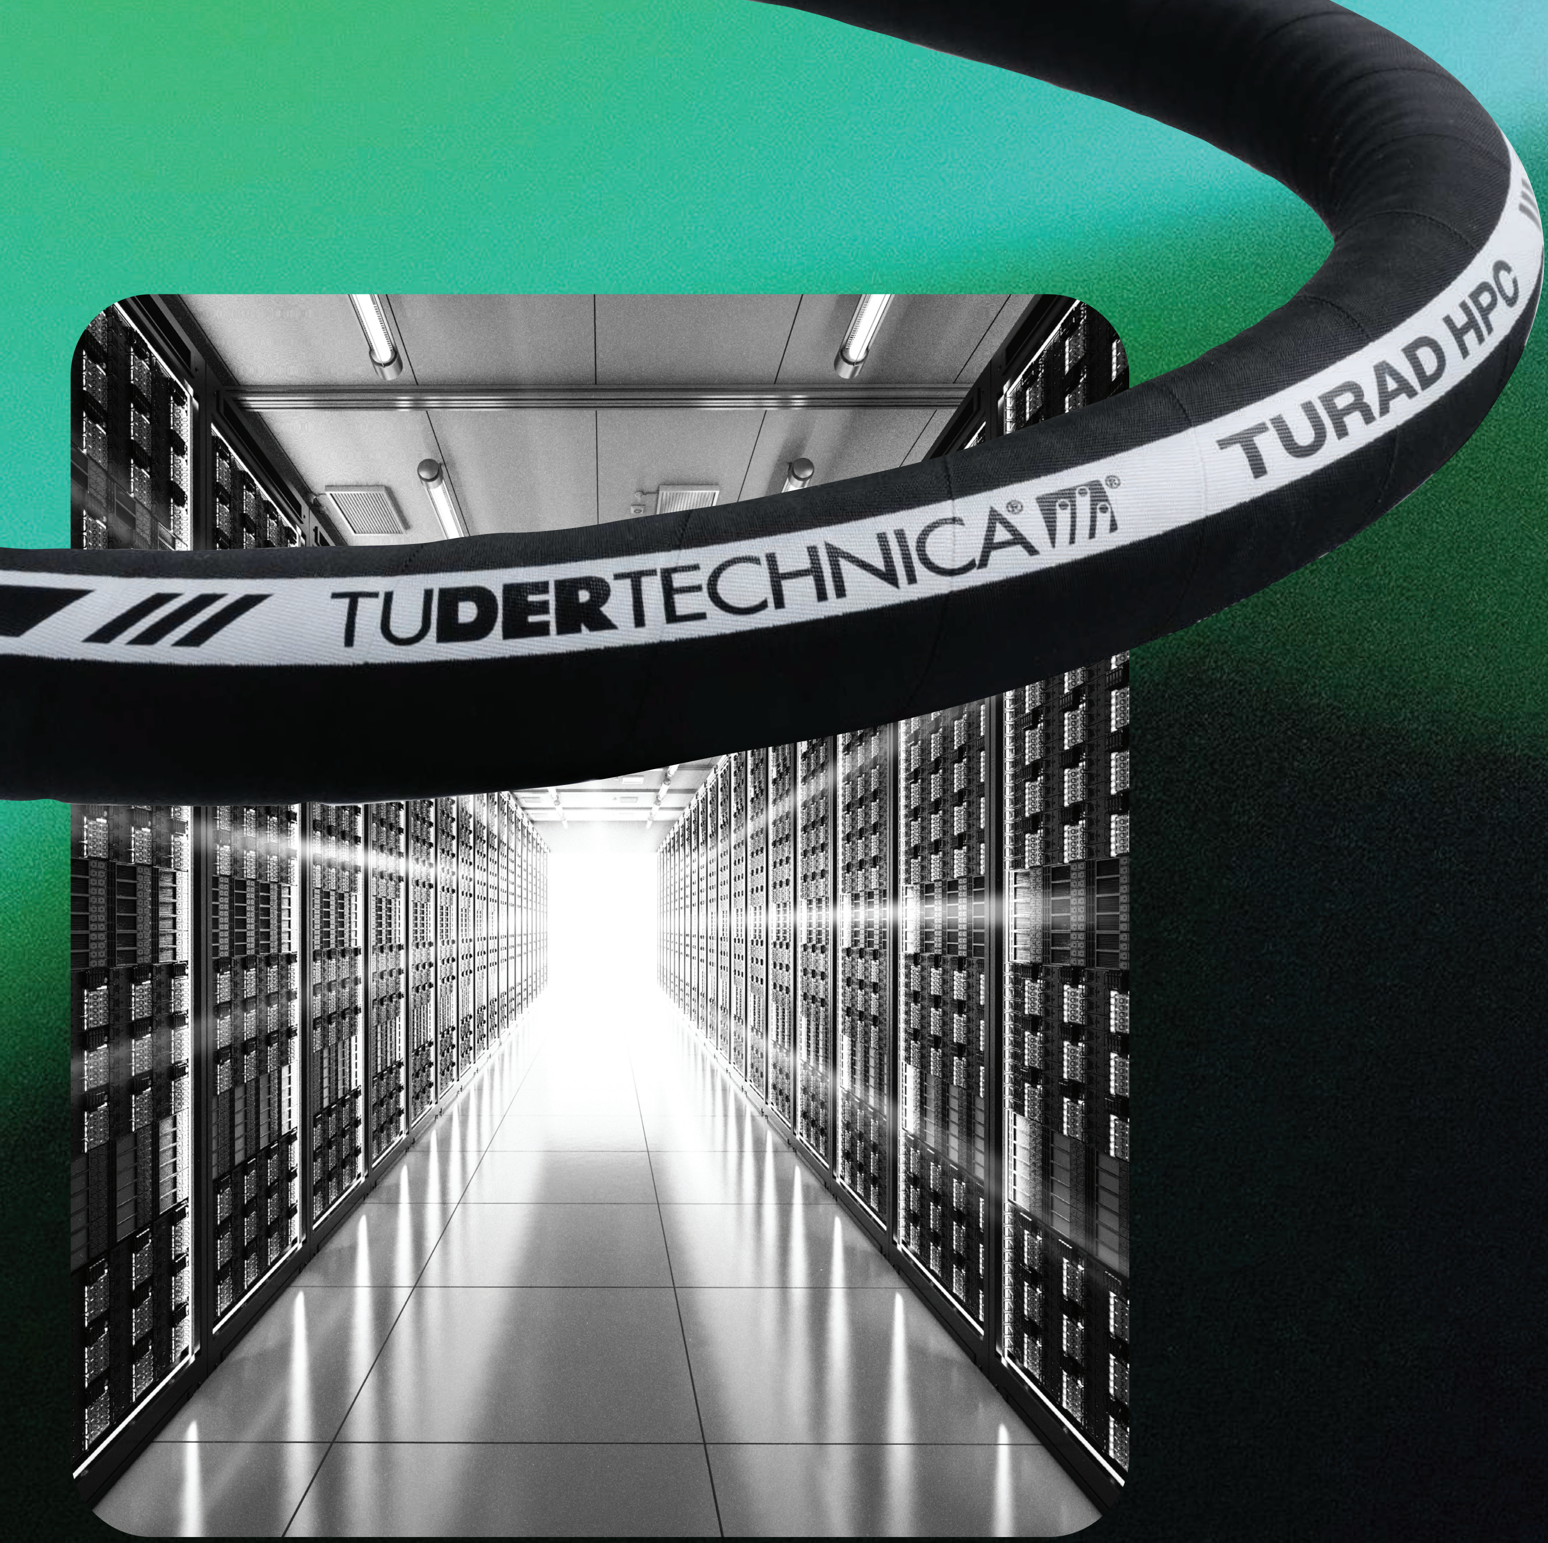


# HPC SERIES

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

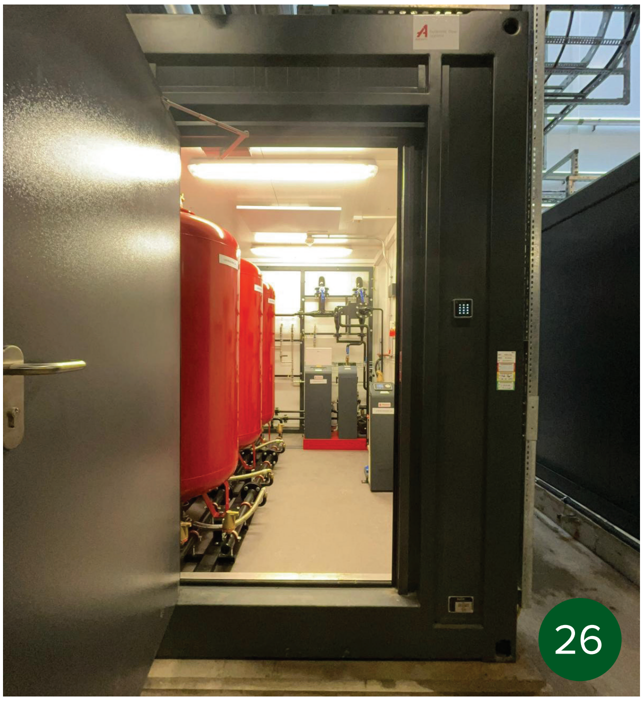





DCS DATACENTRE SOLUTIONS

**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)

# Why EPDs, LCAs and weight matter in data centre infrastructure

Data centre operators are under growing pressure to decarbonise not only the energy that powers their facilities, but also the materials, products and systems embedded into every new build and expansion. In the Nordic region, where carbon reduction is increasingly embedded into procurement, permitting, design and construction expectations, this shift is already visible.

**BY ALAN DURRANT, PANDUIT EMEA**

THE NORDIC COUNTRIES have jointly committed to carbon neutrality in line with the Paris Agreement, and Nordic Energy Research notes that Denmark, Finland, Iceland, Norway and Sweden are using shared clean-energy analysis to track progress against that ambition.

For data centre operations, this matters because the sustainability question is no longer limited to operational energy, Power Usage Effectiveness (PUE), or renewable electricity procurement. The industry is now moving deeper into embodied carbon: the greenhouse gas emissions associated with materials, manufacturing, transport and installation. The International Energy Agency highlights that data centres and data transmission networks are an important and growing source of energy demand, and that their associated emissions depend heavily on the carbon intensity of the electricity systems to which they are connected. But a low-carbon grid does not remove the need to address the carbon already embodied in the electrical and network infrastructure installed on site.

This is where Life Cycle Assessment, Product Carbon Footprint and Environmental Product Declarations become operationally important. A Life Cycle Assessment, or LCA, identifies the environmental impacts of a product, process or activity across its life cycle, including raw material extraction and processing, transport, manufacturing, distribution, installation, use, maintenance and end of life. A Product Carbon Footprint, or PCF, focuses on one specific environmental impact global warming potential.

Global warming potential, or GWP, is the metric that allows different greenhouse gases to be expressed in a common unit: kilograms of carbon dioxide equivalent, or kg $CO_2$e. Carbon dioxide, methane, nitrous oxide and fluorinated gases have different warming effects and atmospheric lifetimes; GWP provides a common technical language for comparing their impact. For data centre owners, consultants and contractors, that common language is essential. It allows procurement teams to ask not simply whether a product is "sustainable", but what its cradle-to-gate GWP is for the manufacturing stages A1-A3: raw material supply, transport to manufacturing, and manufacturing itself. Panduit's LCA guidance specifically identifies A1-A3 as the consistent zone to use for high-level reporting, embodied carbon and cradle-to-gate comparison.

An Environmental Product Declaration, or EPD, makes this data usable. An EPD presents quantified environmental information based on an LCA and is developed using recognised standards such as ISO 14025, ISO 21930 and EN 15804. Panduit's EPDs promote transparency, are valid for five years, and are not intended for unsupported comparative claims; nevertheless, they give customers, specifiers and contractors a structured way to understand environmental impacts and support Scope 3 reporting.

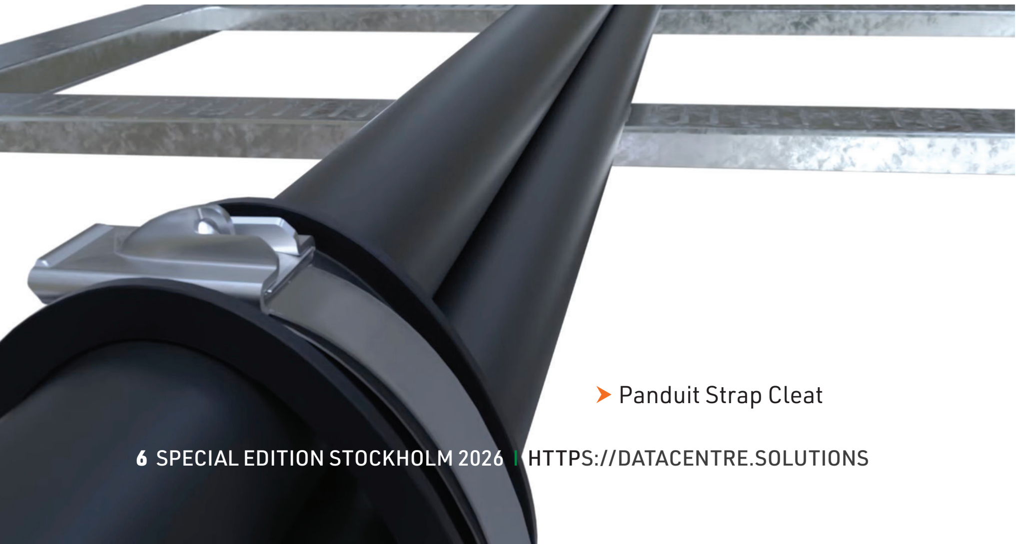

➤ Panduit Strap Cleat

Scope 3 is the critical connection for data centre operators. Manufacturers' Scope 1 and 2 emissions matter, but for a data centre customer, purchased infrastructure normally sits within its Scope 3 inventory. Panduit's response material for data centre requirements makes this point clearly: customers increasingly need product-level GWP data, expressed as kg $CO_2$e per product, and they are looking to manufacturers to provide EPDs for product-stage A1-A3 information. In practice, this means a data centre project can multiply the declared GWP of a cable cleat, jack module, cable or other component by installed quantities to build a more accurate picture of the infrastructures embodied carbon.

The Nordic market is helping accelerate this discipline. Panduit's customer demand for EPDs is being driven first by Norway, followed by Sweden, Denmark and Finland, with the UK, Ireland and Germany following. The data centre sector is a primary driver for EPD requirements in the cleat market. This mirrors the wider Nordic policy environment: Nordic Energy Research reports that the region's journey to climate neutrality is progressing, while also noting that further effort is needed to reach targets in stated timeframes. MLT stainless steel cleat ties show why product design, not just operational efficiency, is becoming part of carbon strategy. Cable cleats are safety-critical components in electrical infrastructure.

Their role is to restrain power cables during short-circuit events, when electromagnetic forces can act within milliseconds and cause major mechanical stress before short circuit protection systems operate. The latest stainless steel strap cleats are designed to provide the required restraint performance while reducing size, weight, installation time, complexity and storage burden compared with traditional rigid cleats.

The carbon significance lies in material mass. Panduit's EPD states that around 95% of the carbon footprint in the cleat is driven by raw material weight, with manufacturing process factors contributing much less by comparison. These cleat solutions have a 10:1 weight advantage over traditional cleats, creating an inherent carbon advantage before other supply-chain factors are considered. In other words, the most effective way to reduce product carbon impact is often not to offset it after the event, but to engineer less material into the product while maintaining performance.

This is a crucial point for data centre operations. Electrical rooms, grey-space infrastructure, modular power systems and high-capacity distribution routes contain large quantities of metal, plastics, cabling and containment.

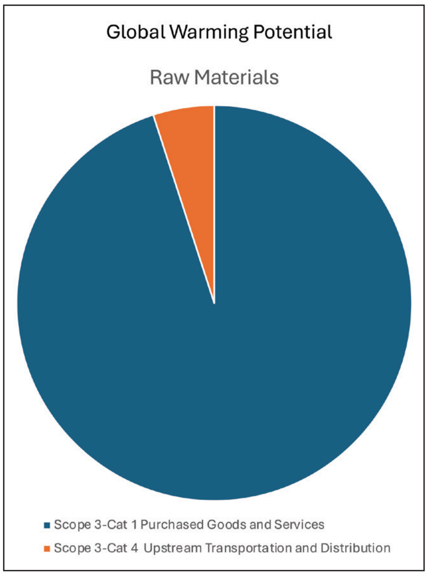


Reducing manufactured product weight is crucial to GWP.

If a product can perform the same certified safety function with significantly less material, its cradle-to-gate GWP can fall because less raw material is extracted, processed, transported and manufactured. This offers broader sustainability benefits that demonstrate the same principle in network cabling: smaller diameter Category 6A cables and patch leads use less material and enable more efficient shipping, with more reels per pallet and improved warehouse, jobsite and transport efficiency.

For MLT steel cleat ties, the operational benefits reinforce the carbon logic. The latest MLT Cleats require substantially less packaging that traditional rigid cleats. Additional benefits include faster installation, reduced part complexity, substantially cheaper, less wasted product and less packaging and fewer sharp edges after installation. These are not only contractor productivity benefits; they reduce logistics intensity, packaging requirements, site congestion and the risk of over-ordering and residual waste. Panduit's wider environmental stewardship programme gives this product-level work an organisational context. The company's stated strategy focuses on carbon footprint reduction, circular economy support, sustainable product innovation and waste elimination. On a corporate level its goals including a 50% reduction in Scope 1 and 2 carbon footprint by 2030, a 10% reduction in energy intensity by 2025, increased renewable or recycled packaging content, reduced single-use plastic packaging and more than 90% diversion of waste from landfill. Through these strategies Panduit has also achieved an EcoVadis Silver Medal sustainability rating.

For data centre operators, the practical message is clear. EPDs and LCAs are not paperwork exercises; they are decision tools. LCA and PCF allow design teams to identify carbon hotspots, compare material strategies, and specify products that reduce embodied carbon at scale. In Nordic-led projects, where low-carbon expectations are increasingly becoming normal practice, cradle-to-gate transparency gives operators a way to move from aspiration to measurement.

The next phase of sustainable data centre infrastructure will be defined by this combination: renewable power for operations, efficient cooling and power distribution, and lower-carbon materials in the build. Panduit's MLT stainless steel cleat ties demonstrate how a safety-critical component can contribute to that transition. By reducing product weight while preserving performance, and by supporting EPD-based cradle-to-gate reporting, Panduit gives data centre operators and contractors a more factual route to reducing carbon impact, one component, one kilogram of material, and one kg $CO_2$e calculation at a time.

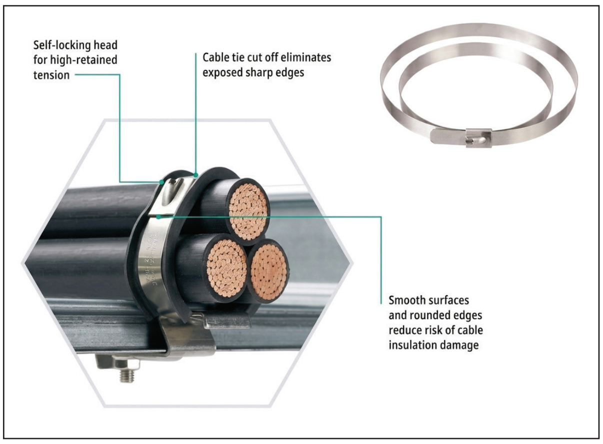


Panduit MLT Strap Cleat.

# 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**<br>Power<br>Use<br>Effectiveness | How efficiently a facility delivers power to IT equipment | Facility Energy & Operating Cost | Facility Engineers<br>Operators | Cooling Architecture Selection<br>Electrical Design<br>OPEX Reduction<br>Efficiency Retrofits | *IT Productivity*<br>*Power Scarcity*<br>*Stranded Capacity*<br>*Site Value* |
| **IUE**<br>Infrastructure<br>Utilisation<br>Efficiency | The amount of installed IT capacity doing useful work | IT Capital Utilisation | IT Leadership<br>Operations Teams | Server Consolidation<br>Workload Placement<br>Refresh Cycles<br>CAPEX Improvements | *Power Availability*<br>*Permitting Limits*<br>*Site Expansion Potential* |
| **PCE**<br>Power<br>Compute<br>Effectiveness | The amount of useful compute delivered per unit of scarce power | Power-to-Compute Value | Executives<br>Investors<br>Utilities<br>Policymakers | Brownfield vs. Greenfield Investment<br>Stranded Power Recovery<br>Site Prioritisation<br>CAPEX Deployment | *Operational or Financial Value Creation Under Power-Strained Conditions* |
| **ROIP**<br>Return On<br>Invested<br>Power | The return generated per megawatt secured | Financial Return on Power | Boards<br>Investors<br>Utilities | Capital Allocation<br>Power Contracting Strategy<br>Long-term Infrastructure Planning | *Technical Causes of Underperformance without PCE Input* |
| EFFICIENCY METRICS OPTIMISE SYSTEMS THAT ALREADY EXIST<br>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.

# 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.

# the lifeblood of every modern data centre

aalberts
mission critical technologies

quality, customisable solutions across the whole data centre building, helping you win time, reduce costs and seamlessly integrate into systems

scan to explore


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.



# Is colocation the catalyst for digital transformations?

Colocation has always been the go-to for businesses that need a safe and dependable home for their data. But now, it's not just about servers. It's about creating a vibrant space where companies can thrive, push boundaries, and conquer the constantly evolving digital world.

**BY FRANCESCO FONTANA, ENTERPRISE MARKETING & ALLIANCES DIRECTOR AT ARUBA S.P.A**

IT'S NOT just about the physical space anymore. It's about building a thriving colocation ecosystem. That means super-fast connections, a variety of cloud providers, and a network of businesses working together to spark new ideas and breakthroughs. As tech keeps evolving, colocation providers who get behind AI, sustainability, and community building will be the ones paving the way forward.

## An increasing risk of data deluge

Data centres are the engines that drive our digital world, handling the ever-growing need for computing power. But with society producing more and more data – from social media, streaming services, and online shopping to 5G networks, cloud computing, and even the Internet of Things (IoT) – the pressure is on for data


centres to keep up and keep delivering. That being said, companies are ready to embrace data and become even more data-driven, with plans to grow and revamp their IT systems. But here's the catch: many of these companies have grand plans for digital transformation, but they're forgetting one crucial ingredient – their data centre.

Something needs to change to ensure their capacity and connections can keep up in the long run and to keep everything secure and running smoothly.

## Could colocation be the solution?

If companies want to grow, they need rock-solid, dependable IT systems that can handle their everyday workloads and then some. If their systems can't keep up, they risk major headaches like downtime, which can mean lost productivity, lost money, and even unhappy customers.

To tackle this issue, some companies have looked to the cloud as their storage saviour, hoping to power new applications. While the cloud's impressive computing and operational space can certainly fit the bill for big enterprises, there are a few factors to consider for the rest, and affordability is a primary consideration.

Cloud storage costs can quickly accumulate and become difficult to manage due to unpredictable pricing models and data egress fees. As a result, many companies have brought some of their workloads back home to on-premise data centres, using a hybrid approach for their IT setup. This way, they get the best of both worlds: on-site data storage and processing power combined with the benefits of the cloud.

But keep in mind that this setup is complex and needs significant networking muscle to keep the data flowing smoothly. And that's where data centre colocation steps in as a real game-changer for big companies.

Think of colocation as renting an apartment for your data. Companies pay for space in a data centre, which gives them the power, cooling, and a safe place to keep their IT gear instead of having to host it all on-site. But colocation isn't just about 'rack space' anymore. It's become a major player in making digital transformation happen, giving businesses the flexibility, scalability, connections, and security they need to push their digital projects forward. By using these professional

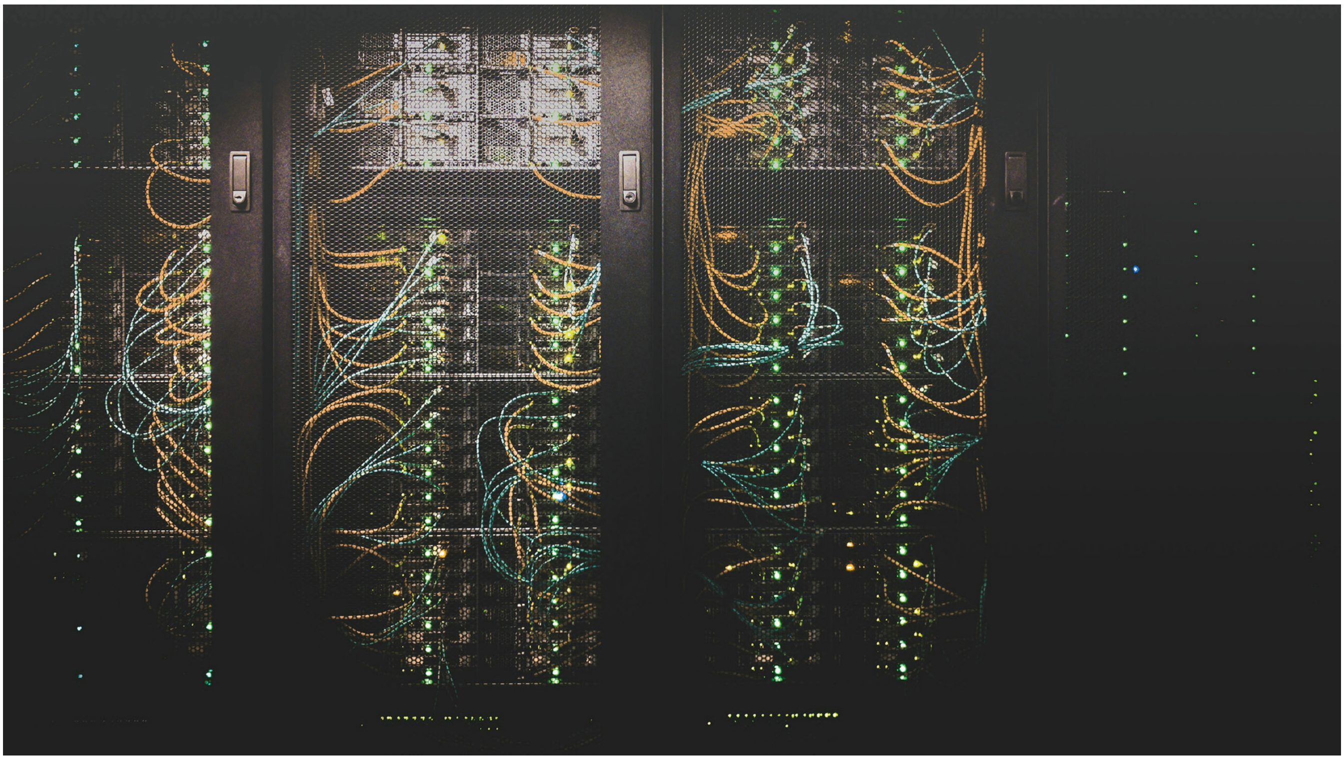

data facilities, companies get a secure and reliable infrastructure without the hassle of building and maintaining their own space.

## How colocation enables digital transformations

Digital transformation requires companies to be quick on their feet and ready to grow. Colocation offers the power to scale IT resources up or down whenever needed. That doesn't just mean adjusting their infrastructure as things change; it's also crucial for supporting new applications, diving deeper into the cloud, and handling all that growing data.

In addition, colocation data centres create ecosystems of connectivity. Companies get direct links to multiple internet providers, cloud services, and other businesses. This makes digital transformation much smoother because it's easier to build powerful IT systems, tap into cloud services without a hitch, and run applications lightning-fast, no matter where you are. And by taking away the hassle of running data centres, colocation simplifies IT management, freeing up valuable time and resources. That means business leaders can focus on other digital projects that actually help the company grow.

While location isn't usually a problem for connectivity, it can be a real game-changer. The closer your data is to you, the faster you can get to it. It is for this reason we see colocation data centres operating near major cloud regions. Countries like Italy and Spain, for example, benefit from strong connectivity to North African and Middle Eastern markets, as well as North and West Europe. We've seen real results from this: lower latency (less lag time) and better connections between companies. That translates to a smoother user experience and even lower broadband costs for those power-hungry applications.

And there's more. Colocation also supports growing worldwide ESG requirements. Going green is a significant priority already, but with all the pressure on businesses to balance digital transformation with being carbon neutral, colocation becomes crucial. Companies need a future-proof plan. Colocation providers understand that, and they've had sustainability as their top priority for years.

## The future is already here

In our world, where data is at the heart of everything, IT systems need to be lightning-fast, adaptable, and dependable if businesses want to go digital without a hitch. Colocation is the perfect answer, giving companies the power of all those cutting-edge technologies, such as the cloud, without the exorbitant price tag.

As more companies jump on the digital transformation train, the colocation market will keep moving forward, growing steadily. As it grows, expect to see even more services emerge, such as the highly sought-after AI-powered resource management and eco-friendly energy solutions.

In our world, where data is at the heart of everything, IT systems need to be lightning-fast, adaptable, and dependable if businesses want to go digital without a hitch. Colocation is the perfect answer, giving companies the power of all those cutting-edge technologies, such as the cloud, without the exorbitant price tag

# Faster Installs. Smarter Power. Zero Complexity.

Managing and installing a rack power distribution unit (PDU) has never been easier than with the EL2P PDU. **Designed to simplify deployment and take the stress out of power distribution,** this intelligent PDU helps reclaim valuable hours.

### Seamless Integration, No Extra Hardware

Built-in data collectors eliminate clutter and the need for secondary devices

### Advanced Connectivity, Maximum Control

Dual 1GB Ethernet ports, Wi-Fi and Bluetooth allow for real-time power management anywhere

### Plug and Play Integration

Seamless integration with third-party DCIM applications


For more information, visit **panduit.com/el2p**

**PANDUIT™**

# 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.


DATWYLER
SMART MODULAR
DATA CENTRES
SCALABLE. ENERGY EFFICIENT. RELIABLE.

DCA CORPORATE PARTNER

# 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.

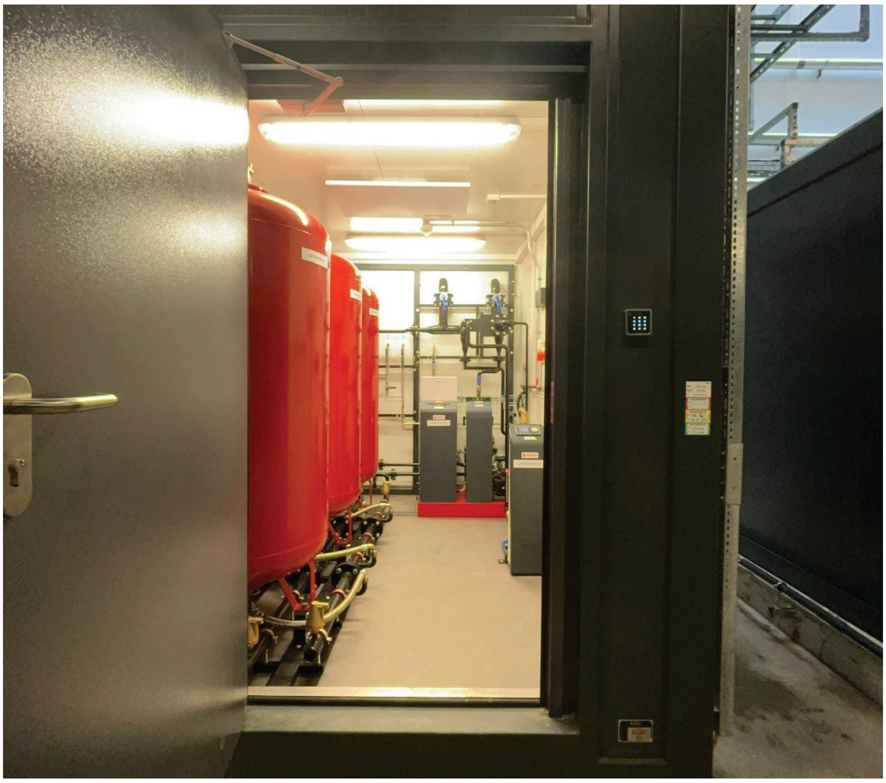

This could manifest in several ways: larger valve sizes in primary plant; different connection strategies in distribution networks, or material and temperature considerations closer to the IT load.

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 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.

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

## 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

# 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

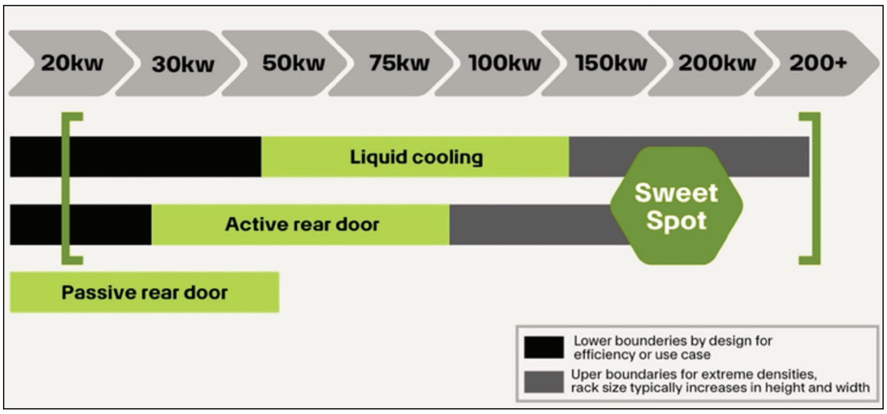


Comparison of Cooling Types

- 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

# STOCKHOLM TECH SHOW 2026

# DATA CEN EXPO 2

## 26-27 MAJ | KISTAMÄSSAN | STOCKHOLM

NTRE
026


## GOLD PARTNERS

katmai · NVIDIA · nextron · Carrier · MERCURY

## EVENT PARTNER

SALBOHEDS · prysmian · akea group · SWECO · AXIS · tietoevry


# 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.


AFL

**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.


25-26 MAY | KISTAMÄSSAN | STOCKHOLM
DATA CENTRE EXPO 2027
SWEDEN'S DATA CENTRES - POWERING THE DIGITAL FUTURE
The journey continues in 2027
Data Centre Expo unites the industry's boldest minds to define, build, and power the future of digital infrastructure.
datacentreexpo.se
stockholmtechshow.se
part of STOCKHOLM TECH SHOW

# AFL at Data Centre World:

## Building the optical layer for AI Scalability

Experience AFL's innovative fibre technologies, enabling next-generation hyperscale infrastructure for AI and cloud. Explore advanced solutions for white space, AI-GPU connectivity, Data Centre Interconnect (DCI), Test-and-Inspection, and Fujikura fusion splicing. See how modular, scalable connectivity, global availability, and customer-focused design simplify and accelerate hyperscale growth.

### On the stand: Fujikura multi-core fibre and hollow-core fibre splicer technology

Discover advanced Multi-Core Fibre and Hollow-Core Fibre fusion splicing technology from the industry's original pioneer. Single and mass-fusion splicers will be on display, including the 100S, 100R, and 45S models. Fujikura splicers are renowned for precision, reliability, and ease of use, whether in cutting-edge research labs or long haul DCI links. Every machine is rigorously engineered to deliver consistent, high-quality splices across a wide range of applications.

### AI-GPU architectures: NextGeneration data centre fabric

AI continues to drive innovation and rapid modernization across hyperscale data centres. High-performance GPUs, dense interconnects, and tightly coupled compute clusters place extreme demands on optical systems. At AI scale, inter-GPU communication dominates fabric traffic, requiring throughput and latency to remain predictable under all operating conditions. Meeting these challenges at the physical layer ensures consistent bandwidth delivery and supports clustering architectures that depend on massive east-west communication.

At AI scale, inter-GPU communication dominates fabric traffic. Optical infrastructure must therefore sustain extremely high data flows with minimal signal degradation. Under these conditions, fibre alignment precision, connector quality, and low-loss routing directly influence performance. These parameters cannot be adjusted later without structural changes, meaning disciplined fibre network design and structured cabling at installation become essential.


**Key requirements for hyperscale deployments supporting AI and cloud include:**

- High-fibre-count ecosystems that retain optical fidelity under load
- Installation practices focused on repeatability and comprehensive test verification
- Proven reference architectures designed for ultra-dense, high-bandwidth workloads
- Access to AFL's worldwide network of optical fibre specialists trained to global standards
- Structured cabling frameworks that simplify management and maintain predictable routing

AI clusters continue to compress hardware footprints and generate higher localised thermal output. As a result, optical systems need to maintain accessibility while enabling rapid scalability in confined white

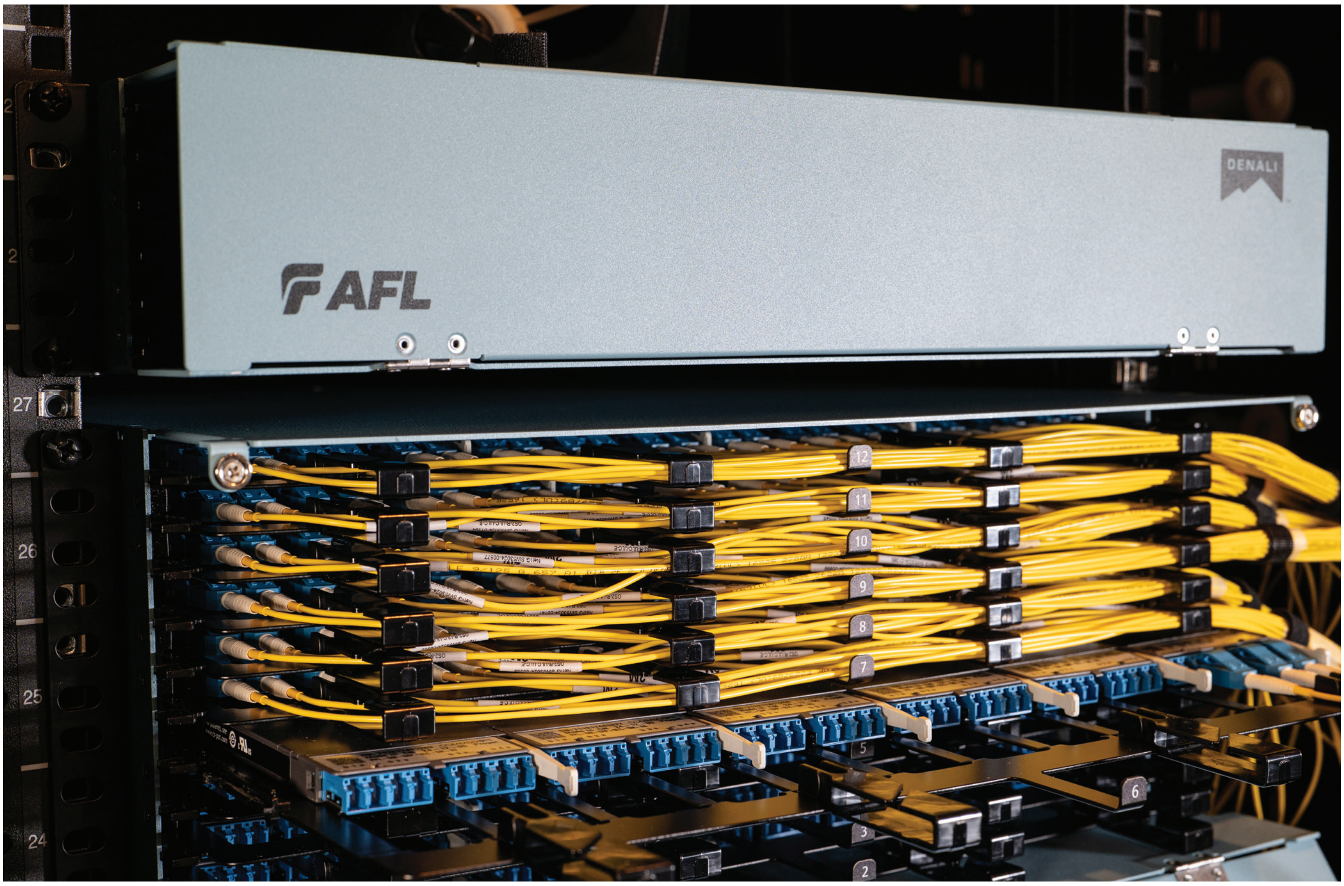


space. This is achieved through modular architectures and precisely routed assemblies, which allow facility operators to expand compute clusters seamlessly while maintaining performance and optimising both physical density and airflow efficiency.

AFL's globally available, standardised connectivity frameworks support replicating designs at scale, while early engagement helps data centre teams plan efficiently, anticipate future upgrades, and align infrastructure investments with long-term growth.

## Precision test-and-inspection for high-performance fibre networks

AFL delivers advanced fibre Test-and-Inspection tools to reliably validate overall network performance, support highly efficient commissioning, and simplify ongoing maintenance across evolving hyperscale data centre environments.

Instant access to comprehensive diagnostics for every fibre link enables engineers to quickly identify potential faults, ensure compliance with industry standards, and optimize signal integrity. By combining precision measurement, traceable reporting, and intuitive operation, AFL's Test-and-Inspection solutions help data centre teams reduce downtime, improve operational efficiency, and confidently scale fiber networks.

## One optical partner for the entire lifecycle

AFL provides end-to-end capability across the optical lifecycle, covering manufacturing, connectivity design, testing, inspection, and long-term support. This comprehensive approach eliminates discrepancies between components, ensuring uniform, high-performance results across distributed networks.

Total lifecycle alignment improves operational consistency and reduces overhead. Backbone fibre, interconnects, and structured cabling are benchmarked against the same high-performance related requirements.

This uniformity minimises integration risk, supporting standardised network monitoring and maintenance practices. Structured documentation further streamlines infrastructure audits and compliance activities across multi-tenant environments.

## DCI for AI and cloud: Scalable, reliable optical connectivity

For AI and cloud deployments, AFL delivers high-density, small-diameter DCI cabling, balancing latency, redundancy, and optical performance across multi-site and multi-region environments. AFL offers cable and connectorised options, plus the hardware to connect at each end, with high-density structured pathways, plug-and-play modules, and traceable documentation to simplify installation and audits.

AI and cloud ecosystems increasingly rely on flexible inter-campus and inter-region connectivity, where DCI networks serve not only as bandwidth backbones but also as strategic links between compute clusters and storage availability zones. To ensure these long fibre routes perform reliably, networks must combine low-loss splicing, temperature stability, and an optimised balance of latency and redundancy.

To meet these demands, AFL's precision optical manufacturing sustains predictable, low-loss attenuation across extended paths. Additionally, integrated certification methodologies confirm project requirements before commissioning, reducing both deployment time and future maintenance windows.

AFL solutions support integrated fibre monitoring and remote diagnostics, equipping operators with actionable insight into optical health metrics such as signal degradation trends or splice loss anomalies. This continuous visibility enables pre-emptive maintenance, supports sustained uptime, and, when combined with standardised backbone designs and rigorous documentation, allows operators to expand campus footprints quickly and confidently.

## White space: Structuring the dynamic core of the data centre

Within the white space, infrastructure must remain adaptable to accommodate continual evolution in compute and connectivity. Because GPU-driven designs frequently alter rack layouts, cable topologies, and airflow patterns, AFL provides connectivity architectures that preserve stability while allowing high-frequency reconfiguration.

By employing structured patching, clear labelling, and organised routing paths, AFL reduces operational risk and simplifies troubleshooting. In addition, modular cassettes and high-density enclosures allow technicians to install or remove links efficiently, simplifying migrations to next-generation optics.

Beyond installation, AFL incorporates verification and cleaning procedures directly into operations. By integrating fibre/port inspection and automated cleaning tools into standard workflows, quality is sustained well beyond the commissioning stage. This visibility and predictability in the white space translate directly into reduced downtime and faster provisioning as clusters expand.

## Density: Managing high-fibre environments efficiently

Next-generation facilities must handle unprecedented fibre volume while maintaining manageability and cooling efficiency. To support these demands, AFL's high-performance optical solutions enable rapid scalability:

- Structured pathways and rack routing maintain airflow and physical order
- Robust engineering supports higher fibre aggregates with minimal loss variability
- Modular element design enables flexible expansion across racks and compute pods
- Plug-and-play modules lower installation time and improve termination consistency
- End-to-end documentation links each port to active configuration records for simplified audits

These solutions balance performance with operational agility, keeping fibre organised, traceable, and easy to manage while minimising technician touchpoints (even during large-scale upgrades).

## Accelerating time to first token

Across AI and cloud infrastructure projects, build velocity offers a competitive advantage. By combining preconfigured products with site-specific packaging and onsite validation tools, AFL's integrated delivery model compresses project timelines:

- Pre-engineered kits align to site blueprints and arrive ready for install
- Clearly marked packaging minimises interpretation during deployment
- Integrated inspection and loss-testing tools confirm optical performance instantly

This efficient approach shortens handover cycles and reduces field resource demand. AFL's global field support network with regional availability enables distributed teams to coordinate while maintaining localised accountability. The result is reliable deployment at the pace that hyperscale programmes require.

## Global consistency for evolving hyperscale data centre builds

AFL's global manufacturing network provides synchronised quality control, streamlined logistics, and inventory proximity to major build regions.

This alignment reduces supply risk while sustaining identical optical characteristics from one production run to the next. AFL's manufacturing consistency gives hyperscale operators confidence that identical cables from any region will deliver equivalent optical results.

Sustainability is now a central factor in largescale infrastructure planning. AFL's manufacturing operations emphasise material efficiency, waste reduction, and optimised shipping strategies that lower carbon footprints without impacting production speed. Combining environmental accountability with technical rigor creates holistic value across the supply chain.

## Customization fFlexibility: Adapting to deployment realities

No two facilities share identical spatial constraints, vendor ecosystems, or installation preferences. AFL addresses this variability through modular optical platforms designed for straightforward customization while retaining global specification alignment.

Flexibility extends across trunk length, connector type, polarity management, and breakout options, ensuring compatibility with both legacy and next-generation transmission equipment.

Each configuration is validated within AFL's quality framework to guarantee consistent optical and mechanical performance. Built-in configurability maintains operational cohesion as facilities evolve, keeping documentation and labelling unified across all regional deployments.

## Building for future capacity

The evolution of data centres toward greater speeds, increased fibre density, and integrated AI-to-cloud interconnection relies fundamentally on robust optical connectivity. AFL delivers hyperscale-ready fibre platforms engineered for long-term scalability, operational reliability, and consistent global performance.

- Comprehensive lifecycle coverage safeguards quality from backbone through white space
- AI and GPU-centric designs accelerate time to compute and streamline complex network builds
- Scalable global manufacturing ensures consistent quality while adaptable configurations tailor deployments to customer requirements

Through these combined strengths, AFL enables builders, operators, and cloud innovators to deliver infrastructure that scales reliably, replicates globally, and performs consistently.

# Panduit's infrastructure zones built for modern data centres – Data Centre World: Stand: 840

Panduit's stand at Data Centre World is its largest to date at DCW and has been designed as a working showcase of modern data centre infrastructure, reflecting the realities faced by today's data centre owners, operators and designers

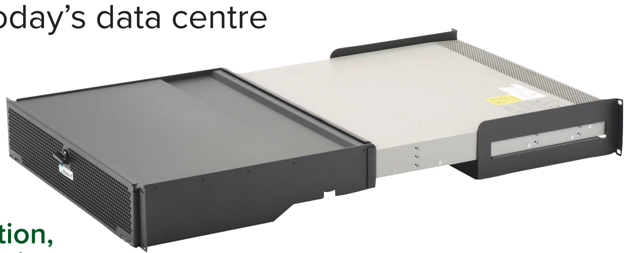

**BY PANDUIT**

RATHER than focusing on individual products, the stand is arranged into zones of interest, each addressing a specific challenge: safety, power and cooling density, structured connectivity, and operational efficiency.

High-Density Cooling and the FlexFusion™ Direct-to-Chip Cabinet
At the heart of the stand is the FlexFusion™ Direct-to-Chip (DTC) Cooling Cabinet, a centrepiece that reflects one of the most pressing issues for data centre buyers: how to cool increasingly powerful AI and HPC servers without sacrificing space or reliability.

Panduit's DTC solution integrates a reinforced cabinet with rack-mounted liquid manifolds designed to deliver coolant directly to server chips. The system supports multiple manifolds and maximised PDUs capacity within a standard cabinet footprint, eliminating the need for rear "caboose" extensions and freeing valuable white space.

Leak-free quick connectors, automatic air exhaust valves and quick-swap components reduce installation risk and simplify maintenance which are key concerns for operators considering liquid cooling for the first time.

From a buyer's perspective, the value lies not only in thermal performance, but in deployment confidence. The cabinet's movable e-rails, bonded steel frame and high static load rating are engineered to support large, heavy AI servers while maintaining accessibility and airflow.

## Power cable protection, grounding and electrical safety

Another major zone focuses on electrical safety and resilience, areas that are increasingly scrutinised as power levels rise across data centre infrastructure. Panduit's cable cleat solutions demonstrate how short-circuit-rated containment can prevent catastrophic cable movement in fault conditions, protecting both personnel and infrastructure. Panduit's innovative single cleat design can accommodate multiple cable sizes and configurations, massively reduce part counts and waste while accelerate installation, a practical productivity gain for large-scale builds.

Grounding and bonding solutions further reinforce this safety-first approach. Panduit's compression-based grounding systems eliminate the need for exothermic welding, reducing installation risk while delivering reliable, standards compliant earthing, a benefit for both installers and long-term operators.

## Structured connectivity and Gen 7 SAN infrastructure

For buyers planning high-performance storage and AI fabrics, Panduit's Gen 7 SAN Director connectivity zone demonstrates how structured cabling underpins performance, scalability and serviceability.

Gen 7 environments dramatically increase port density, and Panduit's solutions address this with high-density fibre architectures that improve airflow and cable organisation while reducing congestion. Pre-terminated SN® connectivity, port replication panels and structured fibre management enable faster installs and easier moves, adds and changes, critical for environments where downtime is costly .

This zone reinforces a key Panduit message, next-generation electronics only deliver their full value when paired with infrastructure that is designed to match their density, speed and complexity.

## Racks, pathways and physical layer integration

Supporting the core cabinet are examples of Panduit's latest two-post racks, wire basket pathways and fibre runners, demonstrating how power and data can be routed cleanly and safely through a facility. Operators can see how Panduit's physical layer components integrate seamlessly from overhead routing to cabinet entry, supporting both traditional and emerging deployment models. This holistic view is particularly valuable for operators balancing legacy equipment with new technologies such as liquid cooling and higher-density fibre.

## Labelling, identification and operational efficiency

The final zone highlights an often-overlooked aspect of infrastructure: identification and day-to-day operations. Panduit's industrial printers, colour-coded labelling systems as well as colour coded cables demonstrate


flamco

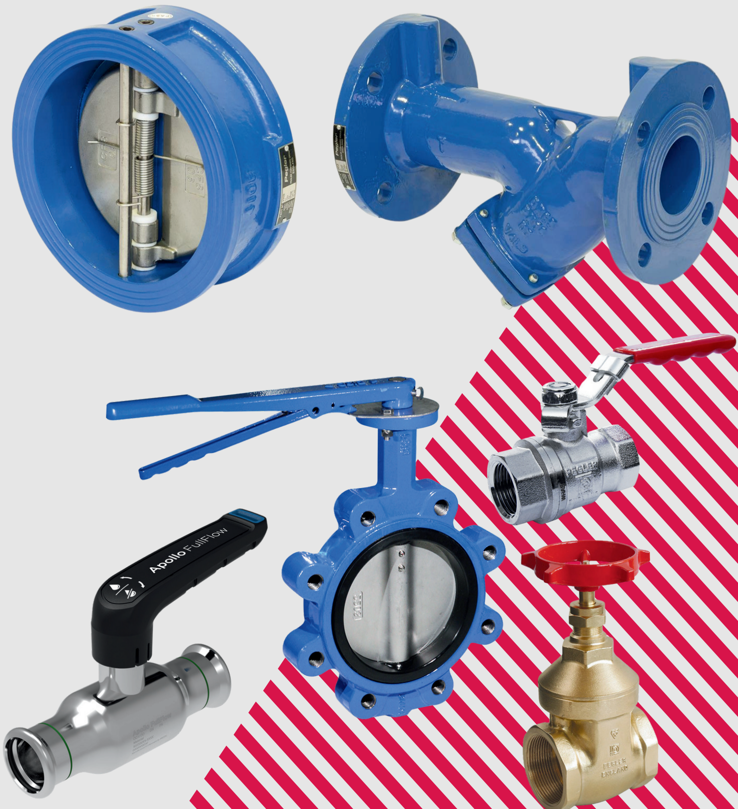
Apollo FullFlow

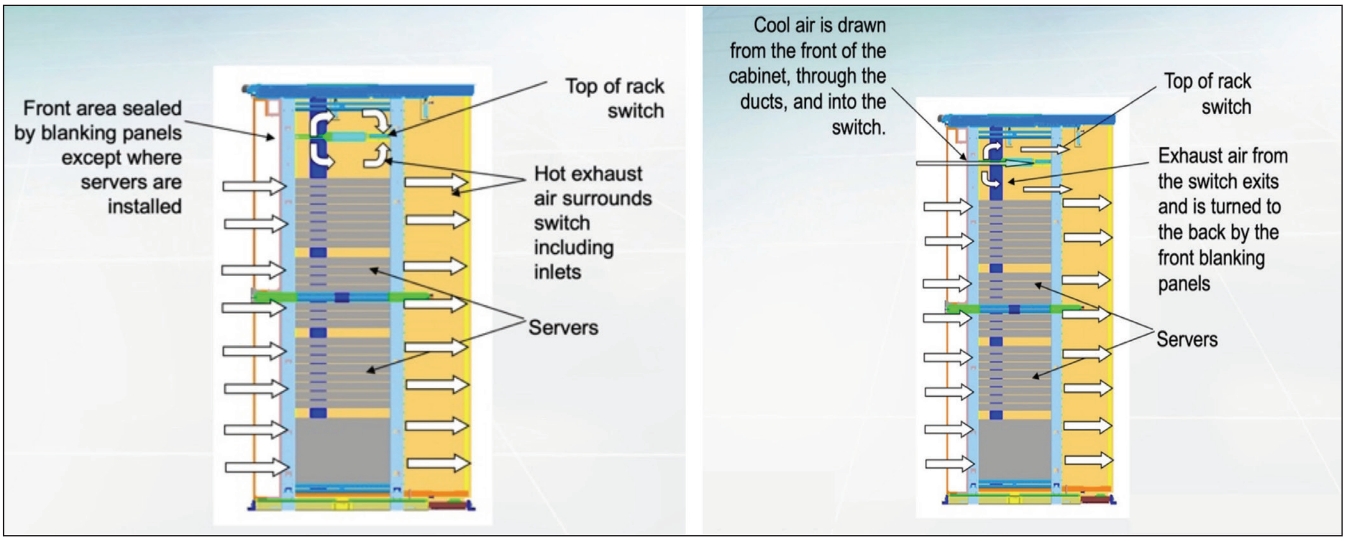


how clear identification improves safety, reduces troubleshooting time and supports compliance throughout the data centre lifecycle.

For infrastructure buyers, this zone underscores that long-term efficiency is shaped as much by visibility and organisation as by hardware performance.

## A complete DCW infrastructure story

Panduit's Data Centre World stand brings these zones together into a single narrative: infrastructure that is safer to install, faster to deploy, and ready for the future.

For data centre buyers navigating rising densities, AI workloads and operational risk, the stand offers a clear message, the physical layer matters mor enow than ever, and when it's engineered as a system, it becomes a strategic advantage.

Come and join us on the stand to understand how Panduit can support your next data centre infrastructure deployment.

# 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.

DATWYLER

# Smart modular data centres

Pre-configured, rapidly deployable solutions that meet your exact requirements—ideal for businesses seeking a reliable, efficient, and future-proof IT foundation for digital transformation.

**BY RALF FISCHINGER, HEAD OF MARKETING & COMMUNICATIONS EUROPE**

DIGITISATION is progressing at a rapid pace – and posing enormous challenges for businesses. With automation, robotics, 5G, the increase of intelligent devices (IoT) and especially the sudden expansion of artificial intelligence (AI), more and more data are having to be analysed, processed and handled. Having enough computing power – anywhere and at any time – has become a central prerequisite for business development.

Powerful access nodes, right at the edge of the network, in the immediate vicinity of the data sources, are important and must nowadays meet the maximum requirements for speed, short latency times, reliability and security. With its Smart Modular Data Centres (SMDC) Datwyler offers innovative, scalable and future-proof all-in-one IT infrastructure solutions, which have been developed especially for edge computing applications. These pre-configured plug-and-play solutions enable companies to process their data efficiently and securely directly at the point of origin.

From collecting and analysing production data to hosting local AI applications: The SMDC solutions from Datwyler provide a versatile way of covering all modern requirements. In addition, these solutions comply with international industry standards for data centres and ensure that each individual component and element can be taken out of service for maintenance or repair without affecting the critical environment or IT processes.

## Real-time data processing at the edge

The key advantage of an SMDC is its short latency times: Data are processed in near to real time thus guaranteeing fast reactions and optimum performance. At the same time, the operation of the data centre remains independent of the network, enabling local aggregation and storage of critical data. This increases reliability and reduces the risk of data loss.

Every Datwyler SMDC is a complete data centre solution that has been tailored to meet the user's exact requirements – fully integrated and

from a single provider. It includes a flexible rack configuration – from a single rack, through multi-rack, to multi-module solutions – with integrated power distribution, an uninterruptible power supply, an efficient cooling system, fire extinguishing system, blanking panels and sealing strip as well as a data centre infrastructure management (DCIM) system that enables complete remote monitoring and control.

The entire infrastructure can be deployed or relocated within a few hours. Thanks to its modular design, every system has the flexibility to adapt to growing demands.


## Energy efficiency

An important aspect of the SMDCs is their energy efficiency. The Datwyler solutions are setting new standards with a power usage effectiveness (PUE) value as low as 1.3. This low PUE value is achieved thanks to a highly efficient cooling system, optimised airflow and intelligent energy distribution. This reduces not only the energy consumption but also the CO2 footprint, thus contributing to a sustainable IT strategy.

The focus is always on security: Integrated access control with both card and keypad input, intelligent PDUs for energy management and real-time status monitoring ensure reliable, secure operation.

In short: Datwyler's Smart Modular Data Centre is an ideal solution for companies who, in an increasingly digital world, need powerful, scalable and reliable IT infrastructures. With the SMDC you are opting for a future-proof solution that will provide you with maximum flexibility and reliability.

**datwyler-itinfra.com**

## About Datwyler IT Infra AG

Datwyler IT Infra is an international company with headquarters in Switzerland and affiliates in Europe, the Middle East and Asia. Datwyler enables organisations around the globe to run their IT and OT infrastructures seamlessly and scale their business with ease.

The well-established company operates as a provider of innovative system solutions, products and services for data centres, fibre networks and intelligent buildings, as well as acting as a subcontractor or general contractor covering the entire value-added chain with tailor-made solutions and outstanding expertise. Datwyler, established in 1915, has a global team of approximately 900 employees.

**datwyler-itinfra.com**

# 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

AVTRON POWER SOLUTIONS

# Combining liquid-cooled & air-cooled load banks for effective data center commissioning

As data centers grow in scale and thermal intensity, traditional commissioning approaches are evolving to reflect modern cooling architectures.

MANY HYPERSCALE and colocation facilities are now starting to utilize both air and liquid cooling strategies to support high-density IT loads. In parallel, commissioning engineers are turning to a hybrid approach – using both liquid-cooled and air-cooled load banks – to properly simulate live operating conditions. This article explores how integrating both types of load banks supports accurate, reliable, and scalable data center commissioning.

## Understanding Load Banks in Data Center Commissioning

Load banks are essential tools for simulating real-world power and cooling demands in a new or upgraded facility. They validate system integrity, redundancy, and thermal behavior before live IT hardware is deployed. Load banks are used to test:

- Electrical infrastructure (UPS systems, PDUs, switchgear, back-up diesel generators)
- Mechanical systems (CRACs, chillers, pumps, liquid cooling loops)
- Controls, failovers, and backup sequences

Historically, air-cooled load banks have been the norm. However, the emergence of liquid cooling in new high-performance data centers requires liquid-cooled load banks (LCLBs) for full validation testing.

## The Role of Air-Cooled Load Banks

Air-cooled load banks discharge heat into the surrounding environment and are ideal for simulating traditional server racks. During commissioning, they:

- Validate airflow management strategies (hot aisle/cold aisle containment)
- Test the hall cooling capacity and distribution
- Offer easy setup and mobility


They are typically used during early commissioning phases or in areas where traditional air cooling remains in use.

## The Role of Liquid-Cooled Load Banks

Liquid-cooled load banks are purpose-built to test liquid-cooled environments, which are increasingly common in high-density compute and AI/ML workloads. LCLBs are connected to the facility's cooling distribution system and:

- Simulate the thermal load of liquid-cooled servers
- Verify pump performance, flow rate, delta-T, and heat rejection capacity
- Stress-test cooling loops and backup modes
- Identify commissioning bottlenecks such as air locks, low flow, or thermal lag

Because liquid-cooled systems discharge heat into the liquid, using air-

| Requirement | Air-Cooled Load Bank | Liquid-Cooled Load Bank |
|---|---|---|
| Simulates air-cooled racks | ✓ | X |
| Simulates liquid-cooled racks | X | ✓ |
| Tests HVAC Systems | ✓ | X |
| Tests CDU/RCU and liquid coops | X | ✓ |
| Verifies total power and heat rejection | ✓ | ✓ |
| Supports phased commissioning | ✓ | ✓ |

cooled load banks alone will not reveal how the liquid cooling infrastructure behaves under real load.

## A Combined Approach

In modern facilities where both liquid and air-cooling technology is utilized, deploying both liquid and air-cooled load banks offers the most complete test environment in the commissioning phase.

By placing air-cooled load banks in legacy rack zones and liquid-cooled load banks in high-density zones, commissioning teams can validate true mixed-mode operations. This reflects the real diversity of cooling demands seen in production. The data center operators can have complete confidence both types of cooling infrastructures will perform effectively in real world operation.

## Networking & Common Control

A hybrid load bank network enables centralized control of both air-cooled and liquid-cooled load banks, providing commissioning teams with a unified interface to simulate and monitor diverse thermal and electrical loads across the facility. By integrating both types of load banks into a single control platform – whether via a wired or wireless control system – operators can coordinate test sequences, ramp loads in tandem, log performance data, and respond to system behavior in real time.

This centralized approach is particularly valuable in hybrid-cooled data centers, where maintaining synchronized testing across traditional and high-density liquid-cooled zones is critical. Hybrid network control enhances efficiency, reduces human error, and supports more comprehensive data capture during commissioning and integrated systems testing (IST).

Utilizing a common control system for both air-cooled and liquid-cooled load banks significantly enhances commissioning efficiency by streamlining operations, reducing manual intervention, and enabling synchronized load testing across diverse cooling zones. A unified interface allows operators to monitor, control, and sequence all load banks from a single workstation—eliminating the complexity of managing separate systems.
This integration supports faster test execution, coordinated ramp-up/down cycles, and comprehensive data logging for compliance and reporting.

This level of integration – particularly when it comes to liquid cooling – requires deep expertise and robust engineering capabilities that not all providers are equipped to offer. Suppliers may offer only air-cooled solutions or provide liquid-cooled units with limited control capabilities, forcing commissioning teams to piece together systems from multiple vendors. This fragmented approach increases setup time, complicates troubleshooting, and can introduce delays – especially problematic under the tight project timelines typical of hyperscale and colocation buildouts.

In contrast, leading data centers prioritize reliability, repeatability, and expert support—making a single-vendor solution that offers fully networked, hybrid-capable load banks highly valuable. A one-stop shop for both air and liquid-cooled load banks ensures not only interoperability, but also access to experienced engineering teams, field support, and commissioning services that can adapt to evolving site requirements. This is critical for minimizing risk and meeting go-live deadlines with confidence.

## Conclusion

As data centers become more complex, load bank strategies must evolve accordingly. Air-cooled load banks continue to play a vital role, but they cannot replicate the demands of liquid cooling systems. A hybrid commissioning plan that integrates both technologies is essential for verifying high-density infrastructure, maximizing uptime, and ensuring the long-term success of mission-critical environments.

Contact us on:
E: customercare@avtronpower.com
T: +44 1780 920100

https://avtronpower.com

### FURTHER READING

- https://avtronpower.com/load-banks/liquid-cooled-load-banks/
- https://avtronpower.com/load-banks/
- https://avtronpower.com/resources/articles/why-use-load-banks-in-data-centers/
- https://avtronpower.com/resources/articles/why-network-load-banks/

# Faster Installs. Smarter Power. Zero Complexity.

Managing and installing a rack power distribution unit (PDU) has never been easier than with the EL2P PDU. **Designed to simplify deployment and take the stress out of power distribution,** this intelligent PDU helps reclaim valuable hours.

### Seamless Integration, No Extra Hardware

Built-in data collectors eliminate clutter and the need for secondary devices

### Advanced Connectivity, Maximum Control

Dual 1GB Ethernet ports, Wi-Fi and Bluetooth allow for real-time power management anywhere

### Plug and Play Integration

Seamless integration with third-party DCIM applications


For more information, visit
**panduit.com/el2p**

PANDUIT™