DCS
DATACENTRE SOLUTIONS
DEVELOPING DIGITAL INFRASTRUCTURE IN A HYBRID WORLD

ISSUE VII 2025

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

HTTPS://DATACENTRE.SOLUTIONS

VIEWPOINT

BY PHIL ALSOP
EDITOR

Liquid gold?

Liquid cooling is rapidly emerging as one of the most important technologies shaping the future of data centres. As demand for high-performance computing, artificial intelligence, and hyperscale cloud infrastructure accelerates, the limitations of traditional air cooling are becoming increasingly apparent. Air-based systems require vast amounts of energy to circulate cooled air, often leading to inefficiencies as heat densities climb within racks and server components. Liquid cooling, by contrast, leverages the superior thermal conductivity of fluids to remove heat directly from processors, memory, and other critical components, allowing for greater efficiency, higher performance, and more sustainable operations.

One of the key drivers behind the adoption of liquid cooling is the growing thermal output of next-generation chips. AI training models, graphics processing units, and advanced CPUs generate heat loads that air cooling alone struggles to dissipate. Liquid cooling solutions, including direct-to-chip cold plates and immersion cooling, can handle much higher thermal densities, enabling data centres to support cutting-edge workloads without excessive energy overhead. This is particularly critical as the industry seeks to balance performance gains with environmental responsibilities, since liquid cooling can significantly reduce reliance on power-hungry chillers and fans.

Another advantage lies in space efficiency. Air-cooled facilities require extensive aisle spacing, raised floors, and airflow management, which limit density and scalability. With liquid cooling, heat is removed at the source, allowing servers to be packed more tightly and floor space to be optimised. This density improvement not only reduces the footprint of data centres but also opens the possibility for deploying high-performance infrastructure in locations where real estate is constrained. In addition, the waste heat captured by liquid systems can be more easily reused for district heating or industrial processes, creating circular energy benefits that further align with sustainability goals.

The future trajectory of liquid cooling will likely involve hybrid environments in which air and liquid solutions coexist, gradually shifting toward broader immersion systems as the technology matures and standardises. Challenges remain, particularly around upfront costs, retrofitting legacy facilities, and ensuring compatibility across diverse

hardware. However, with major hyperscalers and colocation providers already investing heavily in liquid cooling research and deployments, economies of scale are expected to drive down costs and accelerate mainstream adoption.

Ultimately, liquid cooling is poised to transform data centre design over the coming decade. By enabling higher compute density, reducing energy consumption, and facilitating heat reuse, it offers a pathway toward more sustainable, efficient, and powerful digital infrastructure. As computational demands continue to surge, liquid cooling will shift from being an experimental option to a foundational requirement of modern data centres.

What is yet to be decided is whether or not the server manufacturers and the data centre owners/operators will reach a consensus when it comes to which type(s) of liquid cooling will win out in the hybrid future, or will we end up with a hybrid, hybrid cooling environment, with not just air and liquid cooling, but also different types of liquid cooling in the same facility?!

CONTENTS

DATACENTRE SOLUTIONS

COVER STORY

DCS DATACENTRE SOLUTIONS

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

Senior B2B Event & Media Executive
Mark Hinds
+44 (0)2476 7189701
mark.hinds@angelbc.com

Sales Manager
Peter Davies
+44 (0)2476 718970
peter.davies@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

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

Circulation & Subscriptions
+44 (0)2476 718970 circ@angelbc.com

Directors
Sukhi Bhadal: Chief Executive Officer
Scott Adams: Chief Technical Officer
E: info@angelbc.com

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

The growing impact of water scarcity on Europe's data centres

Europe's data centres face a critical challenge as water scarcity escalates amid climate change, demanding urgent solutions for sustainable digital infrastructure.

AS Europe grapples with a notable rise in wildfires amid record droughts and extreme heat, the stress on water resources for data centre cooling has emerged as a pressing concern. Traditionally, debates in the industry have concentrated on surging energy needs. However, in today's climate, water scarcity is proving an equally pressing threat to the sustainability and resilience of Europe's swiftly expanding digital infrastructure.

Robert Pritchard, Principal Analyst at the data and analytics firm GlobalData, highlights, "Climate change has produced weather patterns from wildfires to flash floods that are becoming more extreme worldwide. Superheated ground results in less

rain getting absorbed into aquifers and instead running over dry soil, often destroying lives and livelihoods. Climate change also means that the requirement for water to cool the ever-expanding base of data centres is likely to become an issue that needs addressing."

Previously, the infrastructure needs of data centres have prioritised vast energy consumption. Nevertheless, solutions like new nuclear technologies are being explored, while countries such as Singapore and Dublin have put a hold on or stopped new data centre construction to preserve energy for local populations.

Pritchard further states, "Data centre providers are using technology innovations to try to address energy demands… given the explosion of AI." Yet, even with a push towards alternatives to water cooling and reduction in usage, the OECD predicts AI tools will require a staggering 4.2-6.6 billion cubic meters of water by 2027, exceeding Denmark's total annual consumption and nearing half of the UK's.

Navigating grid challenges in the AI era: Decentralised energy solutions for the UK

WITH artificial intelligence (AI) becoming a transformative force across industries, the UK faces increasing pressure to ensure its data centres are equipped to handle rising power demands efficiently. However, concerns surrounding grid connections risk hampering critical investments in AI infrastructure. There is now a call to explore alternative solutions to mitigate these bottlenecks.

Leading global data centre developer, Digital Realty, has highlighted the urgent need for reforming Britain's energy grid and planning systems. These revisions are essential for ensuring reliable power supplies to accommodate newly established facilities. With AI-related tasks escalating workloads for data centres, the focus is intensifying on maintaining site efficiency and uptime.

According to reports, the number of UK data centres is expected to grow by nearly 20%. This boom brings growing pains, particularly from grid connection delays and the National Grid's mounting struggles to meet increasing demand. In this setting, companies like Aggreko suggest that decentralised energy might deliver needed relief.

Aggreko's recent whitepaper, Bridging the Energy Gap for European Data Centres, highlights how on-site power generation can be integral to keeping data centre projects on schedule. Such measures are crucial given the echoing concerns across the sector noted by Billy Durie, Aggreko's Global Sector Head for Data Centres, who emphasises that AI's transformative impact necessitates reliable power. Faced with lengthening waits for grid access, decentralised solutions represent an immediate and critical stopgap, allowing uninterrupted progress and capitalising on substantial AI-driven investment opportunities.

A historical challenge has been acquiring necessary equipment in key markets. But amid AI's rapid ascent, consulting supply chains and employing temporary equipment hire for short to long-term energy needs are recommended strategies. Aggreko has deployed Stage V HVO-fuelled generators and advanced SCADA control systems at a vital UK location, showcasing how standby power solutions can ensure site resilience while cutting emissions.

As project deadlines loom, continued development and commissioning of data centres demand flexibility concerning grid independence.

The evolution of data center semiconductors: Navigating the AI revolution

The semiconductor landscape is evolving rapidly, reshaping cloud and AI infrastructure with transformative technological advancements.

THE BACKBONE of global cloud and AI infrastructure is undergoing a profound transformation, led by the semiconductor industry. It's a market at a critical juncture, with explosive growth in AI and fundamental changes driving this transition.

In 2024, the total addressable market (TAM) for data centre semiconductors reached an impressive $209 billion, covering compute, memory, networking, and power domains. This figure is anticipated to nearly double by 2030, reaching close to $500 billion. AI and high-performance computing (HPC) are the key drivers, with generative AI substantially affecting demand for processors and accelerators.

GPUs remain central to AI infrastructure, with Nvidia commanding a staggering 93% of server GPU revenue by 2024. The Yole Group predicts that GPU revenue will soar from $100 billion in 2024 to $215 billion by 2030.

Despite their sizeable average selling prices, GPUs are essential for AI training and increasingly used for inference.

In this rapidly evolving environment, AI application-specific integrated circuits (AISCs) are gaining traction. Tech giants like Google, Amazon, and Microsoft are heavily investing in domain-specific silicon to enhance performance and lessen reliance on Nvidia. As a result, AI AISC revenue is set to leap to $84.5 billion by 2030.

The evolution doesn't stop at compute. Memory architecture is also advancing rapidly. DDR5 adoption is ongoing, while high bandwidth memory (HBM) enjoys exceptional demand, especially for AI training. Computational express link (CXL) is becoming integral,

addressing memory disaggregation and latency issues in emerging server architectures.

Data centre silicon leadership is witnessing shifts as well. American firms, notably Nvidia, AMD, and Intel, continue to dominate.

However, China is scaling its domestic capabilities via strategic investments and policies, even as export controls impact supply chains, further reinforcing goals for sovereign development within and outside China.

The role of startups and market newcomers is not to be underestimated. Innovators like Groq, Cerebras, and Tenstorrent are reshaping the market, demonstrating how non-traditional solutions can rival established players in terms of cost, performance, and energy efficiency.

Eric Mounier PhD, Chief Analyst, Photonics at Yole Group said, "The data centre semiconductor industry is today investigating many approaches. At Yole Group, we investigated this domain in depth and analysed the innovations. Today's solutions are all about control. AI workloads are reshaping what chips are built, how they're packaged, and where they're manufactured."

Hyperscalers lead global data centre expansion

GLOBAL data centre capex is set to soar, driven by AI adoption and hyperscale investment over the next decade.

According to a recent report by the Dell'Oro Group, global data centre capital expenditure (capex) is set to increase at a compound annual growth rate of 21%.

A significant contribution to this surge is expected from hyperscale cloud service providers, who will account for half of the projected $1.2 trillion global data centre capex by 2029.

Baron Fung, Senior Research Director at Dell'Oro Group, highlights, "GPUs and custom AI accelerators now account for roughly one third of total data centre capex, making them the single largest driver of growth." These remarks emphasise the substantial role these technologies play in current and future data centre expansions. This trend is further supported by robust spending across the infrastructure, which includes racks, compute solutions, storage, networking, and physical facilities.

The hyperscalers, known for leveraging vertically integrated solutions and custom architectures, are at the forefront of optimising data centre performance while reducing computation costs. Both public and private sector investments are contributing to the expansion wave.

To cater to this global demand, hyperscalers and collocation providers are anticipated to introduce over 50 gigawatts of new capacity within five years. Although the forecast predicts a transient deceleration in 2026, investments are poised to sustain growth in the the long run.

Data centre automation set for market growth

Exploring the monumental growth of the data centre automation market, driven by emerging technologies and shifting industry demands.

THE data centre automation market is on an impressive trajectory, expected to swell from its present valuation exceeding $11.4 billion to a remarkable $50.2 billion by 2034, according to a recent report from Global Market Insights, Inc.

This growth is largely attributed to the widespread adoption of cloud services, burgeoning social media platforms, and the explosion of IoT devices across various sectors. As enterprises pivot towards digital storage and cloud-based infrastructures, the demand for efficient data centre operations is reaching unprecedented levels.

Implementing automation within data centres greatly enhances operational efficiency and mitigates human errors, fostering meticulous data management. Technologies such as machine learning (ML), artificial intelligence (AI), and cloud computing are crucial, optimising processes, minimising downtime, and supporting predictive maintenance.

These advancements ensure businesses maintain a competitive edge in a swiftly changing digital world. The mounting focus on cybersecurity intensifies the push towards automation, allowing for real-time threat detection and enhanced data protection. As industries shift towards hybrid and multi-cloud settings, data centre automation solutions are increasingly essential, spurring innovative technology developments.

Government programmes promoting digital infrastructure adoption underscore the relevance of data centre automation as a strategic business necessity. The market is divided into two core components: solutions and services. As of 2024, solutions held a commanding 60% share, powered by automation software aiding in seamless resource allocation and task automation. Meanwhile, the burgeoning services segment supports organisations seeking expert guidance in maintaining automated systems. Deployment choices between on-premises and cloud-based solutions highlight preferences for the latter, which held a 57% market share in 2024. With data security as a priority, cloud providers enhance measures with encryption and multi-factor authentication.

GPUs powering explosive AI data centre growth

AS NATIONS across the globe ramp up investments in AI data centers and cloud computing, the term "Sovereign AI" is gaining popularity. These centres, often referred to as "AI Factories", are creating a surge in demand for highly specialised computing chips, particularly GPUs (Graphics Processing Units). A report by IDTechEx highlights the trajectory of AI chips, forecasting significant growth in the next decade.

GPUs are becoming indispensable, capturing a whopping 82% of the AI chip revenue in 2024. By 2025, their deployment is expected to multiply, dominated by industry leader NVIDIA with its Blackwell GPUs. Close on its heels, AMD competes fiercely with its MI300 and MI350 series, securing substantial deals with major technology companies.

Initially developed in the 1970s for basic 2D graphics rendering, GPUs have undergone significant transformations. The 1990s witnessed a growth in 3D graphics, with AMD and NVIDIA developing technologies that allowed GPUs to harness parallel processing capabilities for broader uses, such as simulations and image processing by the mid-2000s.

The surge of interest in AI in the 2010s, propelled by models like AlexNet and ResNet, further cemented the role of GPUs in training advanced AI models. Modern-day GPUs are tasked with facilitating complex AI operations, ensuring high-speed processing and supporting vast library functions needed for deep learning.

Comprised of thousands of cores, each GPU is designed to execute specific instructions simultaneously across numerous data points. Despite their simpler cache systems compared to CPUs, GPUs enhance throughput efficiency, crucial for tasks involving extensive data calculations.

The future will likely see high-performance GPUs adopt advanced transistor nodes, such as 2nm, a move that promises greater efficiency and density. However, challenges persist, particularly with the considerable costs of ultra-advanced lithography equipment and other hurdles, such as increasing heat production and materials limitations.

While custom ASICs and emerging chip technologies challenge the GPU stronghold, GPUs remain dominant, thanks to technological innovations like die-stitching and chiplet 3D stacking. Such innovations increase transistor counts and improve yield rates, though often at the cost of memory speed.

DCPI Market poised for rapid growth amid AI surge

The global data centre physical infrastructure market is on track for a significant surge, driven by AI-ready capacities.

THE Data Center Physical Infrastructure (DCPI) market is set for robust growth, with a projected 15 percent compound annual growth rate (CAGR) from 2024 to 2029, as detailed in a recent report by Dell'Oro Group. The escalation will see the market reach a significant milestone of $63.1 billion by the close of this period, fuelled by AI-ready capabilities intensifying mid-decade.

The forecast underscores accelerated deployments tailored to support surging computing workloads. Alex Cordovil, Research Director at Dell'Oro Group, emphasises how AI innovations are redefining facility designs, raising densities, enhancing power intelligence, and transitioning liquid cooling from a specialist solution to a standard necessity.

Region-wise, North America is leading the charge, with EMEA and China seeing peaks in 2026 before likely moderation. AI sovereignty and export-oriented policies prominently support the ongoing momentum. Operators globally are optimising utility strategies by integrating on-site generation to mitigate capacity expansion challenges posed by power constraints. This report offers detailed insights into market segments, including UPS, thermal management, and rack distribution among others.

Dell'Oro Group's Data Center Physical Infrastructure 5-Year Forecast report gives an insight into the Data Center Physical Infrastructure market. This covers market sizes and forecasts for uninterruptible power supplies (UPS), thermal management, cabinet power distribution and busway, rack power distribution, IT racks and containment, and software and services.

Allocation of manufacturing revenues by cloud service providers, telco, collocation, and enterprise customer segments is also given.

AI-driven growth transforming data center front-end networks

RECENT FINDINGS from the Dell'Oro Group have illuminated a significant evolution in the data center landscape. The surge in artificial intelligence (AI) back-end networks is necessitating increased capacity in data center switch front-end networks for data collection and processing.

Projections indicate that this AI-induced demand will spark a market expansion in front-end networks, predicted to grow at a compound annual growth rate (CAGR) of over 40% between 2024 and 2029. This growth is poised to offer substantial opportunities for existing technology players and aspiring new entrants.

Sameh Boujelbene, Vice President at Dell'Oro Group, observed, "Over the past few years, data center switch sales in front-end networks for general purpose non-accelerated servers have been largely driven by brownfield deployments and upgrades of the installed base, as capacity expansion slowed significantly AI back-end deployments are, however, breathing new life into this market."

She emphasized that AI developments are catalyzing a shift in front-end network demands, requiring new connectivity channels. This emerging need connects accelerated servers not solely to each other but crucially to the front-end network for data ingestion. The implication here is an increased need for high-speed connections that carry a higher cost.

The report also highlighted that
1. Retail leaders such as Accton, Cisco, and NVIDIA among others, are positioned best to benefit from this market shift.
2. The predicted deployment of almost 90 million 800 Gbps and 1600 Gbps switch ports in front-end networks over the coming five years is underscored by advances such as 51.2 Tbps and 102.4 Tbps chip technologies.
3. In contrast, switch port shipments in back-end networks are expected to more than triple those figures, further signifying the burgeoning scale of AI influences.

As AI networks drive this market's evolution, stakeholders should prepare for a demanding environment where speed and capacity are paramount, offering endless possibilities.

The complex energy challenge of AI: Balancing progress with sustainability

AI's escalating energy consumption threatens sustainable goals, pushing companies towards carbon offsetting and legislative adaptations.

ARTIFICIAL INTELLIGENCE (AI) is driving transformative advancements across various sectors, ranging from healthcare to finance. However, this momentum is accompanied by substantial energy consumption concerns. Data centres, the heart of AI systems, consume massive amounts of power, prompting industry leaders like Meta, Google, and Microsoft to explore sustainable alternatives.

The intricacies of AI models, especially the training of expansive neural networks, demand significant computational resources, contributing to considerable carbon emissions. For instance, the necessary cooling in data centres alone imposes a substantial energy overhead.

This energy requirement is further propelled by AI technologies such as autonomous vehicles and predictive analytics, which necessitate continuous data processing.
Consequently, increased energy usage entails reliance on fossil fuels like coal, oil, and natural gas, leading to heightened emissions levels.

Renewable energy sources, although promising, aren't scaling swiftly enough, leaving many data centres dependent on traditional power sources.
Projections indicate a staggering 160% increase in power demand by 2030, with AI's energy requirements intensifying.

Recently, headline-grabbing moves by companies underscore the challenge. Meta's groundbreaking 20-year nuclear power agreement for its Illinois data centers exemplifies the critical role nuclear energy plays in meeting surging AI electricity demands. Google and Microsoft are similarly investing heavily in nuclear solutions to sustain AI without elevating carbon emissions.

Data centres, essential for AI systems, rank among the world's most energy-intensive facilities. In line with the EU Energy Efficiency Directive, operators now must prove ongoing energy improvements, aligning with broader regulations like the European Climate Law, targeting climate neutrality by 2050. In the UK, schemes such as ESOS Phase 3 mandate additional energy reporting and efficiency standards.

Leading European economies, including the UK, Germany, and France, spearhead efforts to transition to net zero. However, data centres within these nations face substantial hurdles in achieving decarbonisation and sustainability goals. A CFP Energy Survey reveals that while most data centres possess a net-zero strategy, realising these targets remains elusive. For instance:

- UK: 94% adopt net-zero strategies, yet 22% fail to meet decarbonisation targets.
- Germany: 90% hold net-zero strategies, with 30% falling short of objectives.
- France: 86% implement net-zero plans, but 14% are unsuccessful.

This data emphasises a daunting reality: despite adopting ambitious net-zero strategies, genuine decarbonisation remains challenging. As AI demand surges, the imbalance may compel prioritising AI demands over sustaining net-zero goals, widespread globally.
As AI-driven energy consumption escalates, major tech corporations, including Microsoft, pivot to carbon credits and voluntary carbon initiatives to counterbalance environmental impacts. Microsoft's recent accord with Re-Green to offset emissions underscores this necessity, albeit illuminating current technology's inability to fully support sustainable AI growth.

CFP Energy's comprehensive strategy includes sustainable construction, advanced cooling systems, and offering voluntary carbon offsetting services to bridge sustainability goals.
They advise several measures:
1. **Sustainable Construction:**
Incorporating low-embodied-carbon materials in data centre designs to minimise emissions from construction to operation.
2. **Advanced Cooling Systems:**
Employing innovative methods like liquid cooling to enhance energy efficiency while sustaining peak performance.
3. **Voluntary Carbon Offsetting:**
Using verified carbon credits to compensate for unavoidable emissions, although transparency and ethical concerns must be addressed.
4. **Collaboration:**
Partnering with governments, utilities, and technology partners for systemic progress and scalable sustainable AI infrastructure.

The regulatory landscape is swiftly evolving, with mandates like the Digital Operations Resilience Act and Corporate Sustainability Reporting Directive enforcing transparency in emissions and energy efficiency. These policies are not mere administrative hurdles; instead, they offer data centres opportunities to spearhead decarbonisation.

Reliable Data Centres Demand Intelligent Power Management Solutions

DSEG8900

7” Colour Parallel Genset Controller

Reliable power is critical in data centres, where even the smallest disruption can lead to data loss, downtime, and serious financial and reputational damage. With rising demand for cloud services, AI, data processing, and continuous uptime, the pressure on secure, resilient, and efficient power infrastructure has never been greater.

The DSE**G8900** has been developed to meet demanding requirements, delivering advanced power management, control, and paralleling for complex multi-set systems. With flexibility, scalability, and resilience, the controller offers precise control and comprehensive monitoring, giving data centre operators confidence their power systems will perform without compromise.

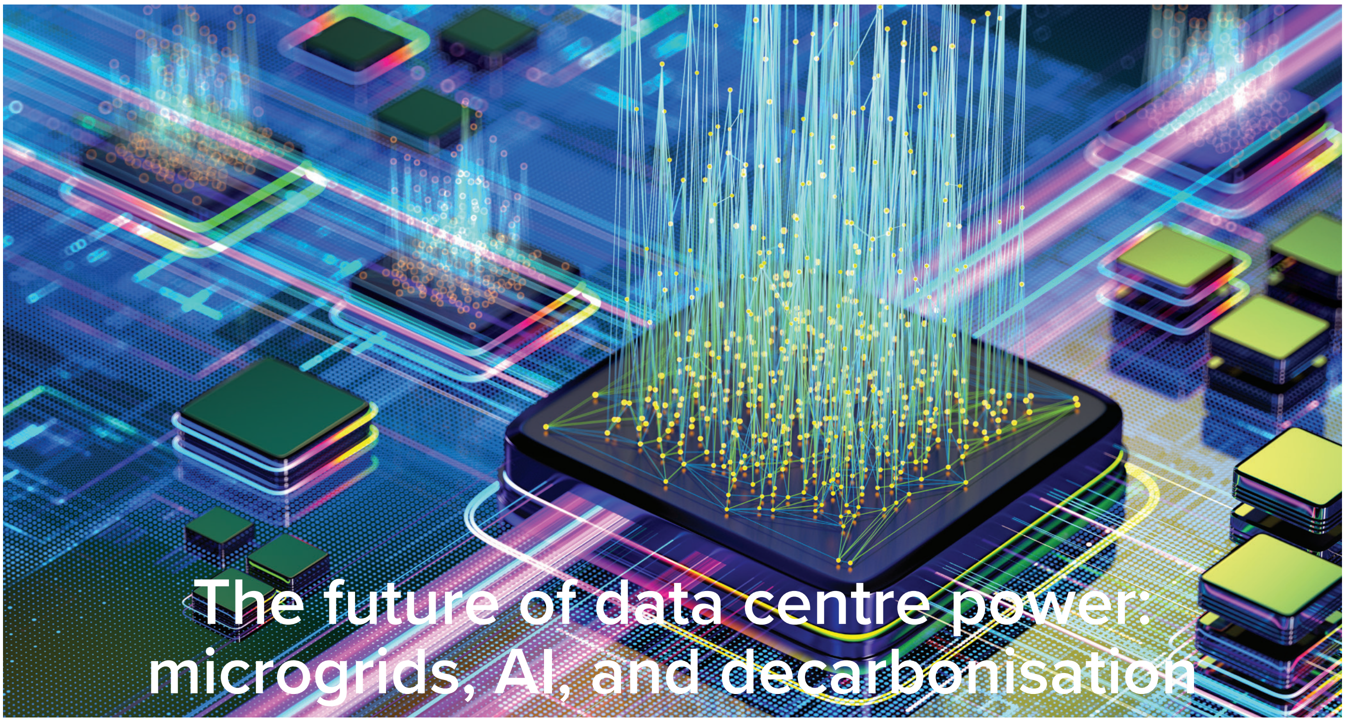

The future of data centre power: microgrids, AI, and decarbonisation

DCS sat down with David Rimmer, Microgrid Solutions Expert at Schneider Electric, to discuss what microgrids mean for the future of data centre infrastructure, how technology is evolving, and why operators may need to rethink their approach to energy procurement and on-site generation.

AS THE data centre industry grapples with unprecedented power demand driven by digitalisation, AI adoption, and sustainability mandates, the conversation increasingly turns to microgrids. These integrated systems promise to improve resilience, optimise energy use, and accelerate decarbonisation — but where do we stand today, and how far can they go?

DCS: *Let's start with the basics — how do you define a microgrid?*

DR: A microgrid is essentially an integrated energy system that consists of interconnected loads and distributed energy resources — generation assets, energy storage, and controllable loads. The important part is that we can control all of this as a single entity.

You can operate in parallel with the grid, or you can switch into what we call 'island mode', running independently from the grid if needed. The ultimate goal is to turn power consumers into prosumers. With solar panels, battery storage, and smart controls, you can choose when to store, sell, produce, or consume electricity. It's about operating your own grid on-site, with the flexibility to stay connected to the main grid or run independently.

DCS: *What are the key benefits of deploying a microgrid in a data centre context?*

DR: There are three big ones: sustainability, cost savings, and resilience.

First, sustainability. If you can produce renewable energy on-site — solar, for instance — that's a major step toward decarbonisation. But it's not just about installing panels and forgetting them. We want to capture and use as much of that energy as possible.

Second, cost. If you can generate energy cheaper on-site than you can buy it from the grid, and you can store and use it strategically — for example, discharging batteries during peak tariff periods — you can significantly cut energy costs.

And third, resilience. On-site generation plus storage gives you backup power capability. Microgrids can supplement grid connections or even reduce your reliance on them, which can be key if you're expanding a campus and the grid connection isn't yet strong enough. Different customers prioritise different drivers, but more often than not, all three play a role in the business case.

DCS: With digitalisation and decarbonisation accelerating simultaneously, does that make microgrids even more important?

DR: Absolutely. But it's not just about having a microgrid — it's about efficiency across the entire chain. Inside the data centre, we need to make operations as efficient as possible. That might mean upgrading UPS systems, installing the most efficient transformers, and thinking about the embodied carbon of equipment choices. Technologies like liquid cooling are becoming increasingly important.

Two data centres with identical server densities can have vastly different energy profiles — with liquid cooling potentially delivering 60–70% energy savings versus air-cooled systems. Once you've optimised energy use, then you can look at decarbonising the supply through microgrids and renewable sourcing.

DCS: *AI is the hot topic right now. How much more pressure is it putting on data centre power infrastructure?*

DR: It's already a major driver. AI and high-performance computing are increasing demand significantly — and quickly.

That makes efficiency even more critical, but it also highlights the need to decouple operations from the grid where possible. In some countries, it's not just generation capacity that's the issue — it's transmission bottlenecks. Microgrids give operators flexibility. By deploying on-site generation and storage, you can reduce your reliance on the grid, supplement capacity where it's constrained, and accelerate project timelines. In what feels like a global race to deploy AI, sites with available, green, and affordable power will win.

DCS: *Is the technology proven, or should operators be cautious about relying on microgrids?*

DR: Most of the technology is mature. Wind and solar are well-established. Backup generators and control systems have been used in data centres for years. What's changing is how we integrate and control these resources.

We can't control the sun or wind, but we can forecast availability using weather data. Add storage to the mix and you can decide when to charge or discharge batteries, when to draw from the grid, and when to sell energy back.

Microgrids are really just about orchestrating these assets in a smarter way. The technology is there — we're just using it differently.

DCS: *Schneider Electric is known for its work in this space. Can you tell us about your microgrid offerings?*

DR: Our flagship tool is EcoStruxure Microgrid Advisor. It's a cloud-based platform that handles energy optimisation.

It uses weather data to forecast on-site generation over the next 24 hours, and it analyses historical data to predict your load profile. That means we know what we can generate, what we'll consume, and when — allowing us to optimise.
If we have excess, we can store it or sell it back to the grid. If we have a shortfall, we can choose when to buy from the grid, based on cost or availability.
And then we have Microgrid Operation, our on-site control system. It makes real-time decisions — starting a genset if needed, switching to battery power, or engaging other assets. Together, these tools combine cloud-based forecasting with local, rule-based controls.

DCS: *Could microgrids be modular, like modular data centres?*

DR: To a degree, yes. You can pre-configure much of the software and control logic off-site. But you still need to connect and commission generation and storage assets on-site. What's interesting is that microgrids, like modular data centres, give you location flexibility. Some AI workloads don't need to be in London or Frankfurt — they can be sited closer to renewable resources.

By co-locating data centres with generation assets — say, a wind farm — and using a microgrid for balancing, you can reduce transmission requirements and even support the local grid.

DCS: *Microgrids also open opportunities for waste heat reuse. Is that realistic?*

DR: Absolutely. If we're moving data centres closer to where the energy is produced, let's also make sure we're near someone who can use the waste heat — whether it's a district heating network, a school, or an industrial process.

And it's not just server heat. Generation assets — CHP plants, fuel cells, biomethane units — also produce heat that can be captured and reused. This can help decarbonise local heat demand while improving the economics of the project.

DCS: *Looking ahead, will microgrids become standard for most data centres?*

DR: You could argue that most data centres already have the building blocks of a microgrid — they've got on-site generation for backup. What's changing is how we use those assets. Instead of just sitting idle until there's a grid outage, we can run them more intelligently, integrate renewables, and add storage. I see microgrids as enabling growth — removing grid constraints as a bottleneck and allowing projects to move forward faster.

DCS: *What about the edge — smaller data centres in remote or distributed locations?*

DR: Microgrids could be even more valuable there. Traditional grids were designed to supply big cities, not rural sites. If you're building edge facilities where grid capacity is limited, a microgrid can supplement that supply and make the site viable.

Smaller facilities also make it easier to balance supply and demand with on-site resources. Whether we end up with a few giant AI campuses or many smaller edge sites — or a mix of both — microgrids will play a role.

DCS: *Given the predicted expansion of data centre capacity, should operators be seriously investigating microgrids right now?*

DR: Yes — because power availability is already a limiting factor in some markets. Microgrids let operators take more control over their energy supply, reduce dependence on central grids, and integrate more renewables.
And they're not a replacement for the grid — they complement it. You still maintain your grid connection, but now you have additional flexibility and resilience. It's about finding the right balance.

Conclusion

Microgrids are no longer just a concept for forward-thinking sustainability teams. As AI drives exponential growth in compute demand and utilities struggle to keep up with grid upgrades, microgrids offer a way to ensure reliable, cost-effective, and lower-carbon power — without waiting years for new grid capacity.

For data centre operators under pressure to expand quickly while meeting net-zero commitments, microgrids might not just be an option — they could be a competitive advantage.

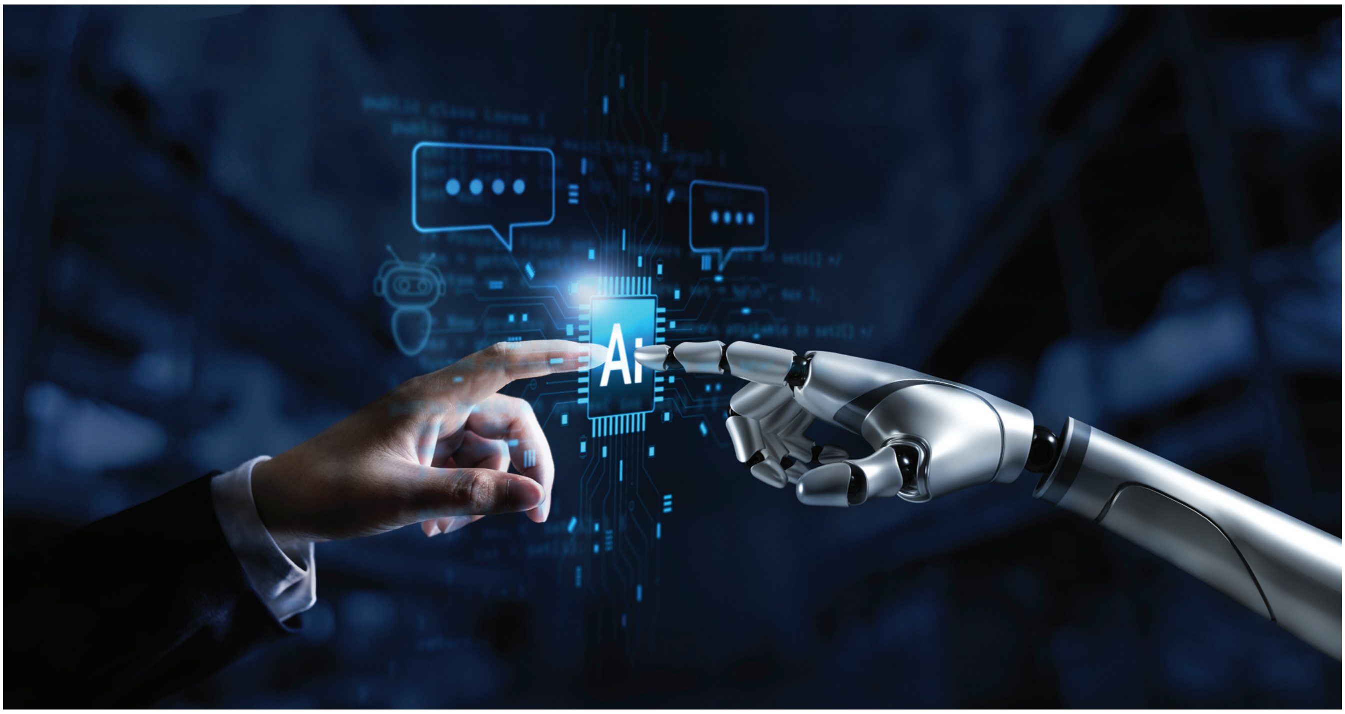

Artificial intelligence can support – but not replace – project managers

At a time when artificial intelligence (AI) is increasingly penetrating many areas of our working lives, the question arises: Can AI replace the project manager? When we ask AI this question, it responds that it can significantly improve project management through enhanced decision-making and increased efficiency. But how far does this support really go, and what are the limits?

BY FRANÇOIS HAYKAL, SENIOR PROJECT CONSULTANT AT BCS, THE SPECIALIST SERVICES PROVIDER TO THE DIGITAL INFRASTRUCTURE INDUSTRY

THERE IS NO DOUBT that the strengths of AI in project management are impressive. In planning and scheduling, for example, it can create automated timelines based on resource availability, task dependencies, and historical data. It can simulate various project scenarios to assess risks and delays and offers recommendations on optimal next steps. In risk management, AI can identify risks more precisely and provide high-quality impact analyses based on vast datasets.

The strengths of AI in project management can be summarised in several areas:

- **Planning and Scheduling:** automated schedules based on resources and task dependencies and simulation of scenarios to assess risks and delays.
- **Risk Management:** precise risk identification and well-founded impact analyses.
- **Resource Management:** optimised allocation of resources based on availability and skill sets.
- §**Project Monitoring:** real-time dashboards for tracking progress and identifying deviations.
- **Document Management:** contract review, risk tagging, and version control.
- **Decision-Making:** recommendations for next steps and efficiency improvements.
- **Quality Management:** error detection, compliance monitoring, and automated change requests.

Facts and figures: AI in cost management

AI seems especially capable in cost management, particularly for data-driven tasks like estimation and benchmarking, quantity take off, invoice and payment verification, forecasting, contract analysis, and reporting. Here, it can fully leverage its analytical strengths and save project managers valuable time. The precision and speed with which AI can process large volumes of data make it an invaluable tool for these aspects of project work.

The limits of AI in project management

Despite its impressive capabilities, AI has its limitations in certain areas. These include in planning where there can be a lack of understanding for nuance and creative conflict resolution and in risk management where it can't evaluate soft risks or undertake holistic decision-making based on "gut feeling."

In the area of resource management, it has an inability to consider personal dynamics such as motivation, fairness, or burnout and when it comes to communication and collaboration it lacks the capacity to build trust, inspire teams, or handle emotional conversations.

The future: AI and humans as a team

The future of project management clearly does not lie in a competition between AI and humans, but in their collaboration. AI is a powerful assistant but not a leader. It is best suited for data-heavy, repeatable tasks, while humans remain superior in leadership, judgment, ethics, negotiation and empathy.

In this context AI can be compared to the autopilot in a modern airplane: it can fly, navigate, and even land under ideal conditions. But a human pilot

is still needed to deal with weather, emergencies, conflicting instructions, and passenger safety. In project management, AI is the autopilot; the project manager is the pilot.

This collaboration leads to faster reporting, better forecasting, and reduced administrative workload – giving project managers more time for leadership, stakeholder engagement, and decision-making.

Responsibility: The insurmountable boundary for AI

In addition, there is one thing that AI will never be able to do in place of the project manager: take responsibility – standing in front of stakeholders or clients to defend a decision or accept blame. Project management is not just about managing tasks; it's about working with people and handling real-world risks and uncertainties, where data may be incomplete and outcomes unpredictable.

Only a human project manager can:
- Make tough decisions when data is unclear.
- Lead under pressure.
- Be the face of a decision.

AI can inform the project manager – but only the project manager can bear the consequences.

Liquid Cooling as a Platform: the missing piece in scalable data centre infrastructure

It's easy to be astonished by how fast AI has progressed. But industry insiders are equally amazed by the pace at which the infrastructure underlying artificial intelligence has developed – and the surge in power demands that comes with it.

BY STUART CRUMP, GLOBAL COMMERCIAL DIRECTOR, LIQUIDSTACK

EVEN the simplest prompt triggers a cascade of computation and data transfer. Every link in that chain consumes electricity. And much of it funneled into powerful, power-hungry NVIDIA GPUs.

Lower-powered alternatives are emerging, but the NVIDIA ecosystem still dominates, dictating the thermal profile of modern data centers. Without advanced cooling, GPUs can't hit peak performance or density.

The International Energy Agency estimates global data center energy consumption will near 1000TWh by 2030, more than doubling the 2024 total. That's a staggering climb, with energy consumption growing 12 percent per year and now accounting for 1.5 percent of all global consumption[1].

But compute is only part of the equation. Every kilowatt powering a chip creates heat. According to ABI Research, 37 percent of the energy used in data centers goes straight to cooling.[2]

A 1MW facility was once a flagship. Today, hyperscalers are designing data centers in the hundreds of megawatts – and NVIDIA is targeting 1MW per rack by 2027. Meanwhile ABI predicts the number of public data centers will quadruple by 2030.

That's not just a growth curve – it's a pressure cooker. And thermal management will define who can scale, who can sustain, and who can lead.

While operators can count on vendors to continually deliver better and more efficient compute, the same can't be said for cooling. Traditional air cooling is commoditized and incapable of handling the heat densities coming with the next wave of AI infrastructure.

The industry is entering a new phase – one where cooling isn't just a backend necessity, but a strategic differentiator.

Liquid cooling is the answer. The Uptime Institute reports that 22 percent of organizations are already using some form of direct liquid cooling (DLC)". Liquid cooling is no longer exotic – but it's still largely custom, especially outside GPU farms and hyperscale environments. That model is not sustainable.

Liquid cooling for all?

What will it take to make liquid the default?
When it comes to servers, storage, or network infrastructure, operators expect easy integration. Cooling should be no different. Whether designing in or retrofitting, liquid cooling must become predictable, repeatable, and scalable.
It's not just about day one. If every cooling system in a data center requires custom design and management, operators can't scale with AI. Cooling must move at the pace of compute.
It must also be easy to service. From hyperscalers supporting global SaaS platforms to enterprise data centers backing up financial services, downtime is unacceptable.

Cooling, platformized

At LiquidStack, we've built our approach around these needs. We started with two-phase liquid immersion – arguably the most demanding form of thermal management. We've since expanded to cover the full spectrum of liquid cooling needs with a major focus on CDUs for DLC systems.

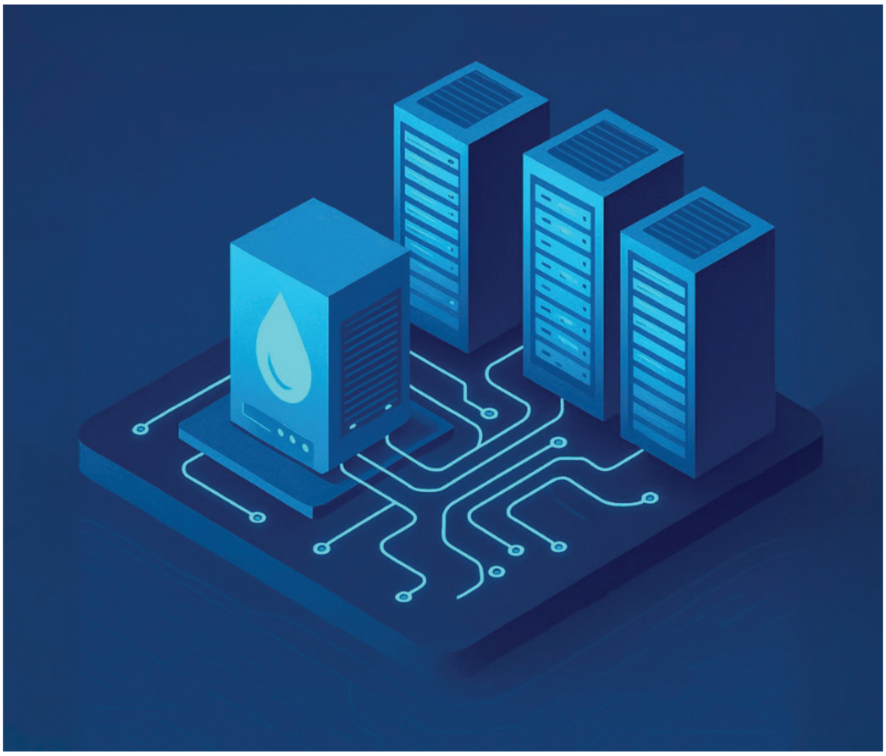

Our latest solution, the GigaModular CDU, is built for scale. It's a single-phase DLC platform that scales from 2.5MW to 10MW, with centralised control and modular pump architecture. Everything is accessible from the front, making service simple and placement flexible.

Operators see a 25 percent saving in capex and floorspace, a critical advantage when deploying rapidly or retrofitting legacy environments. And our "pay-as-you-grow" model helps align capital flows with capacity expansion.

But scale doesn't stop at the rack. We've built resilience into our ecosystem, too – because global operators can't wait on a supply chain. We currently operate two factories in the US and are actively expanding our manufacturing footprint globally. Our global service network ensures consistent SLAs worldwide.

Operators can't afford to slow down – and they can't build past their cooling capacity.

LiquidStack delivers cooling as a platform – scalable, serviceable, and globally deployable – just like the other critical infrastructure in the data center.

FURTHER READING / REFERENCE

- 1. https://www.iea.org/reports/energy-and-ai/executive-summary
- 2. https://www.abiresearch.com/blog/data-center-energy-consumption-forecast?utm_source=chatgpt.com
- 3. https://intelligence.uptimeinstitute.com/resource/uptime-institute-cooling-systems-survey-2024-direct-liquid-cooling

DCA CORPORATE PARTNER

Heat recovery in data centres: turning waste into efficiency

In this opinion piece, wxplore how heat recovery is transforming data centre efficiency. As energy demands rise, harnessing waste heat is becoming a vital strategy. This article examines how it supports sustainability, lowers emissions, and helps shape the future of data centre design.

BY NIGEL MALLON, BUSINESS DEVELOPMENT MANAGER FOR DATA CENTRES AT WEATHERITE

AS DATA CENTRES expand to meet growing global demand, so too does their energy consumption. According to the International Energy Agency, data centres account for approximately 1-1.5% of global electricity use, with cooling systems representing a significant share of that footprint.

In the UK, where the shift toward sustainable digital infrastructure is accelerating, finding innovative ways to enhance efficiency and reduce emissions is critical. One of the most underutilised strategies in this space is heat recovery.

What is heat recovery and why does it matter?

Heat recovery is the process of capturing and reusing heat that would otherwise be expelled as waste.
In the context of data centres, this heat is primarily generated by the IT infrastructure and subsequently expelled by cooling systems working to maintain optimal conditions for high-density servers. Traditionally, this waste heat, or waste energy, would be vented directly to atmosphere.

Rather than venting this thermal energy into the atmosphere, it can be redirected to serve useful purposes such as pre-heating water, supporting HVAC operations in neighbouring buildings, or feeding into district heating networks.

With a typical data centre operating 24/7, the continuous generation of low-grade heat presents a consistent opportunity to improve both energy efficiency and environmental performance. From a sustainability perspective, integrating heat recovery aligns with Net Zero ambitions by reducing the site's overall carbon footprint and enhancing Power Usage Effectiveness (PUE).

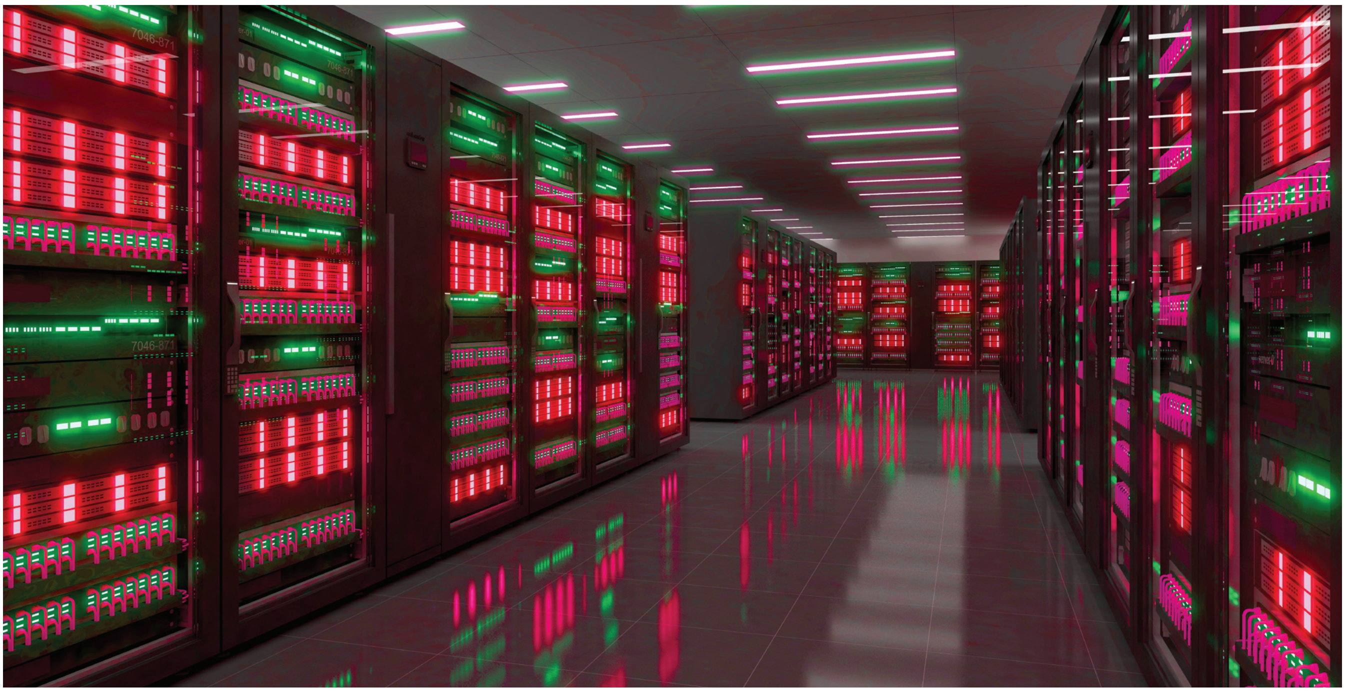

Sector challenges and the opportunity for change

The drive to implement heat recovery comes at a time when data centre operators face mounting pressure to comply with environmental regulations, meet stakeholder ESG commitments, and lower operating costs. Yet, despite the benefits, adoption remains limited due to perceived challenges around integration, system compatibility, and real estate constraints.

Many legacy sites were not designed, or located geographically, with heat reuse in mind, and retrofitting such capabilities can pose logistical and cost-related obstacles. Additionally, engineers must navigate complex temperature differentials, site layouts, and compatibility between HVAC systems and heat recovery equipment.

However, modern HVAC solutions are increasingly being designed with modularity and integration in mind.

Products like Weatherite's Pretec dual cool unit exemplify how innovation in cooling technology can open the door to scalable strategies. The Pretec range combines indirect cooling with integrated heat recovery potential, providing energy-efficient operation even under variable climate conditions. When paired with smart BMS platforms, the unit supports real-time performance monitoring and dynamic adjustment to maximise reuse potential.

Practical applications of heat recovery in data centres

Heat recovery can be applied in several impactful ways within the data centre ecosystem:

- **Domestic Hot Water Pre-Heating:** Captured waste heat can be used to pre-heat water for staff facilities or adjacent buildings, reducing reliance on conventional heating systems.
- **Space Heating for Ancillary Areas:** Administrative offices, loading bays, and common areas can benefit from redirected thermal energy, improving internal comfort while cutting emissions.
- **District Heating Networks:** In urban or campus settings, excess heat from data centres can feed into shared heating systems that serve residential or commercial developments nearby.
- **Support for Humidification Processes:** Facilities using indirect free cooling or requiring specific humidity levels can recycle heat into moisture generation systems, supporting stable environmental control.

Each of these strategies not only helps reduce energy demand but also improves the overall thermal efficiency of the site. According to CIBSE, effective heat recovery can reduce a building's heating demand by 20-30%, depending on usage and design.

Aligning heat recovery with regulatory and sustainability goals

As the UK pushes forward with its Net Zero targets for 2050, data centres are increasingly required to demonstrate energy-efficient design under frameworks such as BREEAM, LEED, and the EU Code of Conduct for Data Centres. Heat recovery directly contributes to higher scores under these schemes by minimising environmental impact and enhancing energy performance.

Additionally, it helps reduce Scope 1 and Scope 2 emissions, a key priority for companies reporting under the Streamlined Energy and Carbon Reporting (SECR) regulations or adhering to Science-Based Targets.

Looking ahead: heat recovery as standard practice

While early adoption of heat recovery in data centres has been sporadic, the combination of environmental necessity, maturing technologies, and economic incentives is shifting the conversation. Forward-thinking organisations are already embedding heat reuse into their design briefs, recognising it as an opportunity to differentiate their operations and add long-term resilience.

Weatherite, with over 50 years of experience in HVAC innovation, continues to support this transition. By working closely with clients in the data centre sector, we're helping to design, manufacture, and implement advanced cooling systems that incorporate heat recovery without compromising operational performance.

Our UK-based manufacturing and engineering expertise ensure that every solution is tailored to the specific thermal and spatial requirements of the site.

Key takeaways

- Heat recovery presents a practical, impactful path to improving data centre efficiency and reducing environmental impact.
- As operators continue to face growing expectations around sustainability and operational transparency, the ability to turn waste heat into a resource rather than a liability will become increasingly vital.
- Through collaborative design, smarter HVAC systems, and a renewed focus on performance-driven engineering, heat recovery can evolve from an afterthought to a cornerstone of future-ready data centre construction.


ENERGY SOLUTIONS FOR ANY TIMELINE
CAT®
POWER FOR THE NOW AND FUTURE
To learn more visit cat.com/datacenter.

Why data centre innovation requires a holistic strategy

Data centres are evolving to stay competitive, maximize performance and safeguard infrastructure. In the race to innovate, investments within these facilities are being made in isolation.

BY MATT POWERS, VICE PRESIDENT OF TECHNOLOGY & SUPPORT SERVICES, WESCO

TOO OFTEN, priority is given to a technology's performance capabilities without considering its relationship with the larger data center technology ecosystem or potential safety, staffing or energy usage implications. These piecemeal, disconnected investments create blind spots or vulnerabilities for data centers, resulting in higher costs, operational inefficiencies and added risk that undercuts performance.

To get the most from new technology investments while minimizing risk, data center operators should consider the capabilities of new technologies along with operational, financial, commercial, security and environmental, health and safety implications.

Seeing the big picture

Data centers today face a set of interconnected challenges:

- Deploying high-performance computing (HPC) to power AI infrastructure, enable large language model (LLM) training and inference at scale, and maximize graphics processing unit (GPU) utilization for the most demanding workloads.
- Using cooling systems that keep up with growing power demands.
- Establishing operational visibility across diverse facility systems to spot and address issues.

A holistic approach that leverages expertise and considers the impact of new technologies from multiple perspectives can help data center teams make more strategic investments in the areas that are critical to their operations today and also sets them up for future improvements as operations evolve. Areas where a holistic approach can make meaningful impacts include:

Aligning HPC to specific data center workloads:

High-performance computing (HPC) platforms are delivered with generalized configurations that are standard defaults for a variety of customers. Unfortunately, AI workloads are anything but generic. To better align HPC platforms with the unique needs of their workloads, some data centers use optimization services. These services precisely match hardware capabilities with actual AI production patterns, such as training, inference at scale or hybrid recommender systems. This fine-tuning turns HPC platforms into true AI factories optimized for throughput, efficiency and scalable production workloads.

Some of the benefits that HPC optimization delivers include:

- Two to ten times improved performance on AI training jobs when systems are properly aligned, depending on baseline configuration and workload type.
- Extended system lifespan through thermal envelope control and smarter workload distribution.
- Faster time-to-insight, especially for LLMs and deep natural language processing (NLP) use cases.

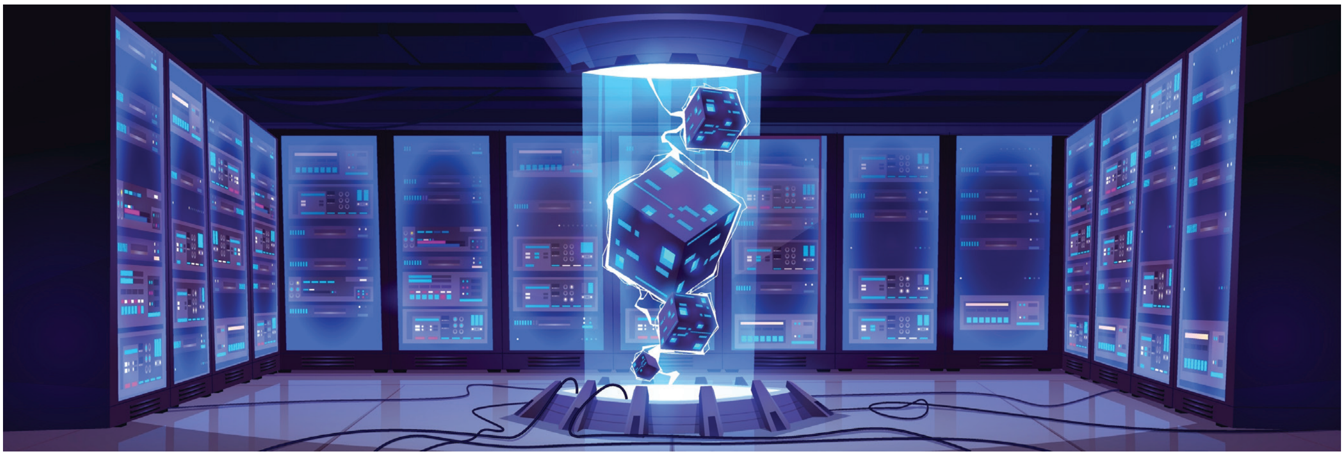

- Improved return on investment per watt, per rack and per GPU.
- The creation of scalable multi-tenant environments with isolated GPU slices that don't degrade Service Level Objectives.
- Visibility and predictability through full-stack observability and regression-aware tuning.

Going "under the hood" of HPC platforms sounds risky but a service provider that's experienced in HPC optimization – and fluent in areas like AI platform telemetry, GPU-level tuning and orchestration stack integration – will keep the process entirely within vendor-supported configurations and tooling and keep platforms within their warranties.

Choosing the appropriate cooling system

Liquid cooling is required in today's high-density data centers, which generate too much heat for traditional air-cooling systems.

Multiple liquid-cooling options are available. To select the right one, data center teams must consider factors like each technology's performance ranges, deployment demands and maintenance requirements.

For example, rear-door heat exchangers provide an efficient and complete liquid cooling solution. When supplied with the correct water temperature and flow as specified by the manufacturer and paired with appropriate containment, they can remove essentially 100% of the heat generated by IT equipment, and the air they discharge from the cabinet is the same as the room's ambient temperature. However, this technology is typically limited to 85kW to 90kW per rack.

Direct-to-chip cooling can support up to 100kW per rack and in some cases go as high as 120kW under optimal conditions. However, this isn't a full-cabinet cooling solution. It only cools the chip, so another solution is needed to cool the rest of the cabinet.

This becomes more important as rack densities continue to increase with innovation and new chipsets potentially reach densities as high as 250kW per rack.

Immersion cooling is another option, but it remains limited in use today because it slows maintenance. Equipment must be lifted from the cooling fluid and then dried before work can be done.
Data center teams should anticipate potential issues associated with their liquid cooling system up front. This includes verifying if a building has the chilled water capacity, supply temperature range, and delta T to support their proposed liquid cooling system. Cooling design directly impacts allowable rack TDP, fan performance, and overall power budgets — which in turn can influence network fabric architecture.

Optimizing operational visibility: Almost as vital as thinking about how different hardware will come together in a data center is deciding how all the data from that hardware will be integrated and used.

> Direct-to-chip cooling can support up to 100kW per rack and in some cases go as high as 120kW under optimal conditions

Rather than monitoring every data stream separately, teams can use a modern data center infrastructure management (DCIM) platform with IT telemetry ingestion to get a single, bird's eye view of all data center operations. A DCIM platform aggregates data from every information technology and operation technology system and device in a data center and helps operators make sense of it, all in a single, integrated experience.

A DCIM platform can deliver useful insights into areas like a data center's energy usage or proactive maintenance needs to optimize facility operations. And with command-and-control capabilities, the platform can ease operators' jobs by allowing them to manage multiple data center functions in one place.

Because the purpose of a DCIM platform is to connect data across disparate technologies, data center teams should choose a platform that is open and vendor agnostic. It should also have built-in integration with facility systems and API connectors for other vendor software.

Empowering the data center ecosystem
Individual technologies don't create competitive data centers. That only happens through smart, coordinated integration of multiple technologies. By thinking holistically about each new investment, data centers can unleash the performance capabilities of new technologies while reducing risk, maximizing investment and better positioning their operations for future changes and new demands.

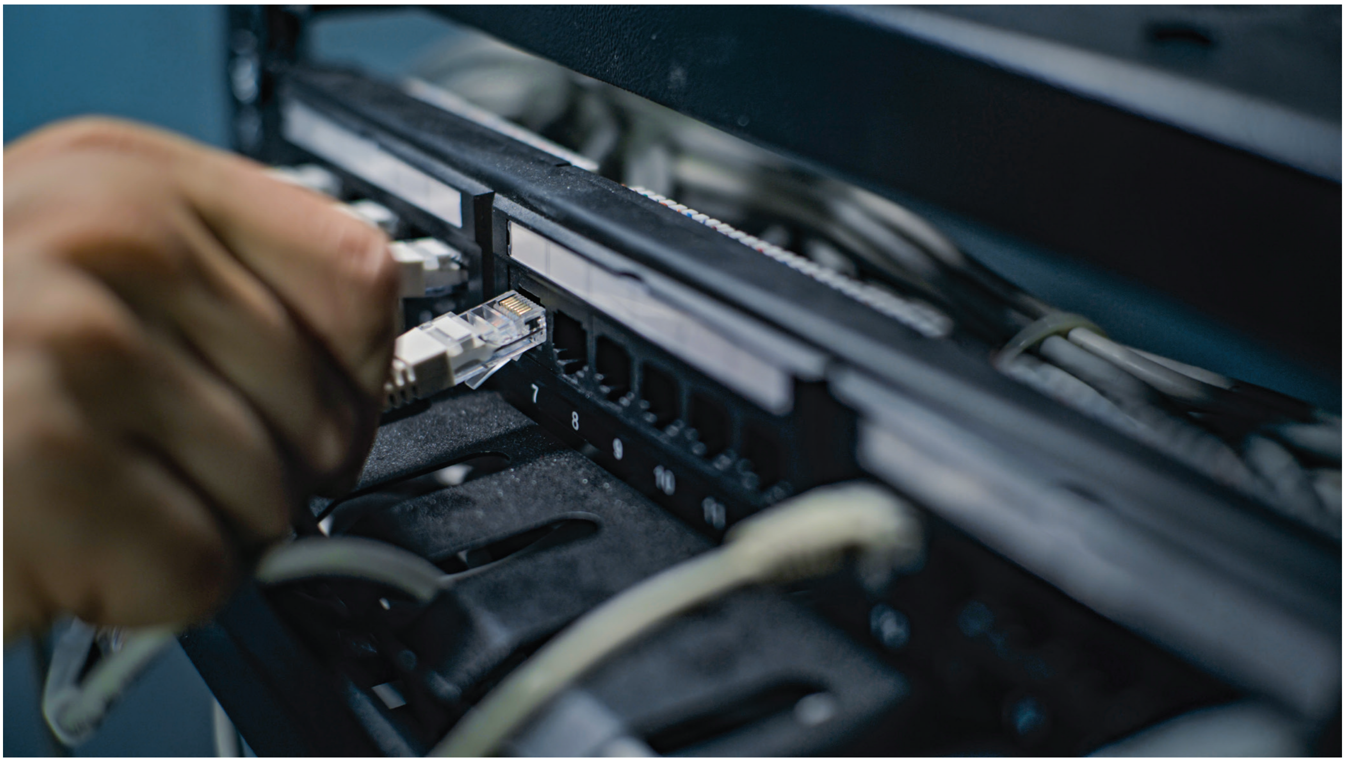

Can the cable cope? Delivering power consistently still relies on cable

It's hardly a secret that one of the biggest challenges facing the developers of data centre facilities is that of power. Most often raised is the issue of grid capacity and the wait time for grid connections.

BY DR RICHARD LEWIS FROM PRYSMIAN HIGHLIGHTS THE VITAL IMPORTANCE OF POWER INFRASTRUCTURE

LESS OFTEN is the spotlight turned on the power cable system. Which, given that the entire operation of the data centre turns on the reliability of this element, is perhaps surprising.

Many facility owners are content to leave the specification and procurement of high voltage power cable to the contractors. As a manufacturer, we'd make the case for at least an understanding of the issues involved in achieving the right result at the best cost.

Cable sizing

Connecting high energy use data centres to the grid may involve cable designed to operate at voltages as high as 132kV. It's essential to design the cable system for the specific power needs of the individual data centre.

Data centres can run on a single supply circuit designed to handle the maximum power requirements, but this can mean that the systems are running at the top end of their performance window, which could unnecessarily stress the system. Consideration should be given to running a double circuit system. In this arrangement, the power load is shared between the two circuits while also building in a degree of redundancy. In the event of a problem with one circuit the other can handle the full load requirements while repairs are completed.

As the complexity of the data centres escalates, specification of the cables should take into account a reasonable expectation of future power demands. This allows the designer to make the right choice of cable right from the start.

No such thing as a standard cable

Establishing the power requirement is the first stage, but cables for this purpose are not "off the shelf" items.

With data centre developments increasingly being planned for centres outside of the traditional distribution network, the length of cable run from facility to grid connection can increase significantly. The installation route will need to be planned accurately and the cable manufactured to specific lengths to minimise the number of joints.

For higher voltage cables the ground conditions will affect the design of the cable itself. Routing that needs to go under a road or river, for example, will need to factor in the derating effect of having to go deeper under the obstacle.

For these reasons, every high voltage power cable will be manufactured to a bespoke design to suit the load and the environment of the route.

This process has been in operation for decades, and each Network Operator will have its own specifications relating to cable design. Each new connection will need to integrate effectively with the existing infrastructure. In certain cases this may mean using specialist joints to connect modern XLPE cable with legacy oil-filled cables.

Only manufacturers approved by the local DNO will be permitted to connect to their network.

The impact of installation

The reliability of the cable system is very dependent on the quality of the installation. Pulling cable through the ducts is itself a skilled operation – the force required to pull the cable must be carefully balanced to its construction. For example, the radius of the bends in the ducts, the height changes over the pull and the conductor material are all factors that need careful consideration to ensure the cable is not overstressed, affecting the long-term performance of the cable.

Cable and accessories are manufactured in a highly controlled environment, with significant quality control processes at every stage of manufacture. Jointing of the accessories, other hand, is down to the skills of the individual operatives working on site, often in less than ideal conditions and under time pressure.

It pays to design the cable system with as few joints as possible, and only to employ jointers who have completed rigorous training programmes specific to the cable and accessory types used in the UK. No cable manufacturer will be able to provide a warranty for complete system unless they also provide the design and jointing services.

Preparing for the worst

A properly designed, expertly installed cable system should perform faultlessly for decades. But, as some recent high-profile circuit failures have demonstrated, there is always the potential for problems.

As a manufacturer, we would always advise on contingency planning from the start. This would include precautions such as:

- Installing an asset monitoring system alongside the cable. Included as part of the initial installation. The additional cost is minimal, and systems offered by Prysmian such as PRY-CAM will provide consistent real-time monitoring. This measures the loading of the cable and identifies any areas of weakness in the cable, joints or terminations, usually in the form of discharge – potentially in advance of any catastrophic failure. This allows for remedial action and circuit outages to be planned in advance. In a worst-case scenario, such monitoring can assist in locating any fault, allowing for fast resolution.
- Holding back-up cable and cable accessories in store. Without this precaution it could take weeks for a factory to be able to schedule the manufacture of replacement products to specific design for repairs.
- Building in redundancy. Most data centres will use battery back-up or independent power generation, and we would also advice using a double circuit for mains power supply for resilience.

Looking to the future

While the energy used in manufacturing and installing power cable will be dwarfed by the operational energy needs of a data centre, there are still significant and positive choices to be made.

The first is to design the cable for optimal efficiency of the whole system. Resistive loss from heavily loaded systems causes energy waste and increases the cost of delivering energy. Then there are manufacturing changes that will have a significant impact on the embodied carbon or associated greenhouse gas emissions, and the types and source of the materials used in the cables themselves.

Prysmian focuses on the use of traceable copper with recycled content and zero-carbon aluminium. The company also has a programme in place to reduce CO2 emissions from our own manufacturing processes. With a target of net zero emissions by 2035, we are already making progress and have reduced our Scope 1 & 2 emissions by 37% compared to a 2019 baseline. Our aim is for 90% reduction by 2035.

We believe that clean, traceable, renewable electrical power is the resource for the future. Building the infrastructure to support that transition is our field of operation and no organisation has more experience of working with the UK transmission and distribution network than Prysmian.

Optimizing UPS total cost of ownership for sustainable data centers

Discover how smart UPS choices can reduce both energy costs and environmental impact in modern data centers

BY MARK MURPHY - EUROPEAN COMMERCIAL DIRECTOR FOR UPS AT LEGRAND

THE DATA CENTER INDUSTRY is facing increasing pressure to accommodate growing capacity demands driven by AI workloads, while also striving to meet sustainability and Net Zero goals. With energy efficiency, operational improvements, and maintenance optimization under intense scrutiny, UPS systems have a critical role to play within the data center, not only in terms of guaranteeing the reliability, resilience, and scalability required to meet growing demands, but also in terms of supporting decarbonization, ESG compliance, and circular economy objectives.

This article examines the crucial factors of a UPS system that can impact Total Cost of Ownership (TCO) and how a lifecycle-based approach to UPS selection can generate long-term savings and sustainability benefits.

TCO: A key metric for sustainable design

TCO combines the initial Capital Expenditures (CapEx), which includes the cost of equipment and installation expenses, with the ongoing, long-term operational expenditure (OpEx), which includes energy consumption, maintenance, cooling, downtime risk, and end-of-life handling.

From a sustainability perspective, choosing a UPS system that is efficient, modular, and environmentally friendly can not only lower electricity bills but also reduce emissions, maintenance overheads, and electronic waste, significantly impacting TCO in the long run.

Silicon carbide: A new benchmark for UPS efficiency

It is widely recognized that higher efficiency in UPS systems directly translates to lower overall energy consumption in data centers. In recent years, significant investments have been made in advancing the technology behind UPS power converters. The introduction of Silicon Carbide (SiC) technology represents a major breakthrough, enabling energy performance levels that exceed 98% efficiency.

SiC has a wide bandgap of 3.26 eV—compared to Silicon's 1.12 eV - allowing it to operate more efficiently at higher voltages, switching frequencies, and temperatures.

These capabilities significantly reduce energy losses, enabling SiC-based UPS systems to deliver exceptional efficiency while minimizing energy waste and long-term operating costs.

Beyond the UPS itself, higher efficiency leads to reduced heat output, which in turn lowers cooling requirements. This contributes to a measurable decrease in the CO_2 emissions associated with electricity consumption. In large-scale data centers, even a modest improvement in UPS efficiency can translate into substantial operational savings and sustainability gains.

Battery management: smarter, longer, cleaner

Batteries are a crucial component of the UPS, essential for ensuring energy continuity. However, they can also be one of the costliest parts of a UPS system, incurring expenses related to monitoring, periodic checks, replacement, and disposal. If batteries are poorly selected and maintained, they can significantly impact the data center's TCO.

A well-chosen UPS system with advanced battery management and control can optimize energy consumption and required capacity (thereby reducing the number of battery blocks and their footprint). When integrated with real-time monitoring and diagnostic systems that communicate with the Battery Management System (BMS), operators gain insights into battery health and performance. Predictive maintenance enhances reliability, ensuring batteries are replaced only when necessary, minimizing both costs and material waste.

This integration is particularly beneficial for lithium-ion batteries, as it extends their lifecycle and aligns with sustainability goals by decreasing the total amount of battery material consumed over time.

Smart grid integration

Modern data centers may encounter risks and costs not directly related to the UPS but which the UPS can help mitigate. As data centers increasingly connect to renewable energy sources and flexible power grids, they may experience costs and downtimes due to grid frequency instability or peak load absorption. Smart grid-ready UPS systems can implement fast frequency response or peak shaving, enabling operators to adapt to fluctuating grid conditions. From a TCO perspective, these features reduce indirect energy costs and improve alignment with renewable energy strategies.

From an environmental standpoint, smart grid integration contributes to greater grid stability and helps accommodate intermittent renewable sources like wind and solar energy, thereby accelerating decarbonization across the energy ecosystem. In short, a smart UPS can act as a bridge between your sustainability goals and your bottom line.

End-of-life and considerations and the circular economy

The lifecycle of a UPS doesn't end at decommissioning, and neither should your sustainability strategy. End-of-life considerations are becoming increasingly crucial as e-waste regulations tighten, and environmental, social, and governance (ESG) scrutiny increases.

Opting for a UPS system with a higher percentage of recyclable materials and enhanced circularity can lower disposal costs and bolster your sustainability reporting, especially when aiming for certifications such as LEED, BREEAM, or ISO 50001. Look for UPS systems that come with a Product Environmental Profile (PEP) to document the percentage of recyclable materials and assess environmental impacts across the entire lifecycle.

Sustainability and TCO go hand in hand

To manage the operational costs of a UPS in a data center effectively, organizations should consider several strategies and actions. These include developing a clear infrastructure lifecycle plan, detailed monitoring, scheduling preventative maintenance, and negotiating a strong after-sales service contract. However, the initial selection of the UPS should be approached as more than just a cost line item.

Choosing a UPS solely based on power capacity, autonomy, and low price—while overlooking other important characteristics and features—can lead to higher operating costs. This decision should be viewed as a long-term investment in resilience and sustainability.

Conversely, selecting a modular UPS system that offers high energy efficiency, quality materials and design, high availability, and intelligent architecture can reduce energy and maintenance costs, lower emissions, support smart grid initiatives, and simplify end-of-life management.

By evaluating UPS technology through the lens of TCO, organizations can align their financial performance with their environmental responsibilities, demonstrating leadership in an era where every kilowatt, euro, and gram of CO_2 matters.

Legrand, a global specialist in electrical and digital building infrastructures, provides a comprehensive portfolio of UPS solutions, ranging from compact single-phase units to high-power modular three-phase systems designed to meet the evolving demands of modern data center environments.

Interested in reading more? Download the full white paper: Optimizing UPS Total Cost of Ownership in Data Centers.

Keor FLEX: High-efficiency, modular, scalable, and sustainable UPS

- Up to 98.4% efficiency in double-conversion mode (99% in ECO mode), reducing energy consumption, cooling costs, and TCO.
- Modular, hot-swappable design supports up to 1.2 MW per system, expandable to 4.8 MW via parallel units, enabling flexible "pay-as-you-grow" scaling.
- Total front access, predictive diagnostics, and smart-grid compatibility simplify maintenance, enhance reliability, and cut operating expenses.
- Incorporates advanced Silicon Carbide (SiC) technology to minimize internal heat generation for improved overall efficiency.
- Recognized for forward-thinking design with a Red Dot Award 2025 - combining usability, aesthetics, and performance.

Learn more about Keor FLEX.

Mark joined Legrand over three years ago as a Key Account Manager, managing the data center solutions portfolio across several strategic accounts. Now serving as Commercial Sales Director for Legrand's UPS business in Europe, he has played a key role in the launch of Keor FLEX—Legrand's next-generation UPS designed to meet the evolving demands of modern data centers.

Is the German Energy Efficiency Act aligned to data centre growth?

Germany has always set the pace when it comes to environmental regulation and the data centre industry is no exception. When the Germany Energy Efficiency Act (EnEfG_ was brought into action at the start of 2024, it brought some of the most demanding energy efficiency requirements the sector has seen.

BY ALEC STEWART, PARTNER, DATA CENTRES, CUNDALL

THE ACT focuses on how data centres are designed, built, and operated, with the aim of cutting emissions and speeding up the shift to renewable energy. But over a year on, there's still a lot of uncertainty about how these rules will work in practice. The industry is questioning if these targets are truly aligned to the future requirements of data centres, particularly those built for AI and liquid cooling, which Germany has the potential to be a hub for.

Looking closer at the Act

The goal of the EnEfG is simple enough on paper. It says to cut down energy usage and find better ways to deal with waste heat. This is mainly by improving how data centres are designed and run. The targets themselves are unambiguous in their language, clearly stating the metrics that must be achieved:

- Data centres that commence operations before 1 July 2026 must achieve a power usage effectiveness (PUE) of 1.5 or lower by 1 July 2027 and 1.3 by 1 July 2030.
- Data centres that commence operations after 1 July 2026 must have an annualised PUE of 1.2.

In addition, the Act mandates the reduction of residual heat produced. Data centres with annual energy consumption exceeding 2.5GWh must reduce heat production to the minimum technically and reasonably feasible – as it is known in German 'zumutbar'. Again, the targets seem clear:

- Data centres that start operating on or after 1 July 2026 must achieve an energy reuse factor (ERF) of at least 10%
- Data centres that start operating on or after 1 July 2027 must achieve a projected ERF of 15%.
- Data centres that start operating on or after 1 July 2028 must achieve a projected ERF of 20%.

However, there are exceptions built into the legislation. If a municipality or energy supplier agrees to develop a heat network within 10 years of the data centre' operational date, they may be exempt from immediate ERF requirements. Similarly, if a local heat network is unwilling to accept the residual heat, the operator only needs to install the necessary infrastructure, typically a heat transfer station. This flexibility indicates that the Act knows some of the technical and financial limitations the legislation could impose on data centre operators. However, given the rise in AI and high-performance computing, which can lower PUE metrics faster than new technologies can raise them, these targets are a concern.

The AI of it all

The timing of the EnEfG Act's release coincided with a huge shift in how data centres are designed. The release of ChatGPT in 2022 and its subsequent introduction into every industry has meant more data centres need to be designed with AI and high-performance computing in mind. As this demand grows, so does the need for computing power. With this comes higher rack densities and therefore, more heat.

It's a trend being driven by hyperscale and enterprise customers, both in Germany and around the world. These days, rack loads of 25kW are standard, and it's not unusual to see peaks hitting 75kW per rack. Data halls are getting bigger too, with many of the latest facilities being built to support 10MW data halls featuring 20x24 rack configurations.

These changes have made it harder for traditional air-cooled systems to keep pace. To account for this, the industry is shifting toward liquid cooling. While this change is needed to deliver the performance that AI workloads demand, it introduces added complexity for operators to meet the Energy Efficiency Act. By lowering the temperatures of chilled water systems, which are necessary for liquid cooling, there is a negative impact on chiller efficiency. In turn, the ambitious targets set by the legislation are harder to achieve.

Is waste heat easier?

In short, no. The Act's requirements for heat reuse present other challenges. Data centres generate large amounts of heat, but finding economically viable uses for it isn't simple.

Across the campus, some heat can be repurposed to heat ancillary or administrative spaces. It could even be used to power equipment like, in the case of generator crankcase heating. But in practice, that only accounts for a small portion of the heat a modern data centre produces, particularly facilities built to handle AI workloads.

Some operators in other countries have been successful by transferring the leftover heat to nearby municipalities or commercial heat networks. This can be a economically beneficial and sustainable solution. However, this only works if the local infrastructure is in place and third parties are willing to invest in making it happen. Even when those pieces are in place, the lower water temperatures that come with liquid cooling can make heat recovery far less effective.

Without viable off-site heat reuse options, operators are left with very few ways to hit the targets set out in the legislation. And even if they plan to rely on one of the exceptions written into the law, those often depend on timelines and decisions made by external parties that are beyond their control.

Managing ambition

As an industry, we need to decarbonise, and the goals of the German Energy Efficiency Act are aligned. They have set clear targets for decarbonisation which is a valuable step toward reaching the industry's goals. However, given the latest developments in AI and the required speed of its adoption, the legislation is creating real technical and operational challenges.

Achieving these low PUE figures and meeting ERF targets whilst still being financially viable is a tightrope. This tightrope needs to be questioned by the industry to see if the Act, in its current form, is well-balanced. Regardless of the Act, if Germany is to achieve its goals, more investment is needed in heat networks. This is not a task just for the government; it will require operators, municipalities, and technology providers to come together to recognise the impact AI will have on how heat can be reused.

Germany has made ambitious targets for the industry, which is a positive thing. In the coming years, more countries will look to Germany's data centre market to see the results the legislation has had. It will set the tone for the next decade, so we need to make sure it is the right one.

Digitisation in the field: construction's next frontier

Intelligent and useful field-centric technologies are reshaping how construction projects are executed and governed. What began as isolated point solutions, RFID tags, a drone here or there, has coalesced into an integrated digital fabric that now layers machine-learning, analytics, computer vision, and predictive AI across scheduling, production control, quality, safety, and equipment management.

BY SHANTHI RAJAN, CEO OF LINARC

EARLY ADOPTERS are already reporting double-digit gains in labor productivity, 30–40 percent reductions in rework, and a measurable drop in lost-time incidents. For executives under pressure to deliver predictable margins and safer job sites, AI-powered field digitization has moved from "nice to have" to strategic imperative.

From static schedules to real-time, AI-driven workflows

The days of printing a two-week look-ahead and hoping it survives the week are numbered. Today, an AI-first field app syncs crews, materials, and inspections in real time; machine-learning engines weigh past productivity and weather to reshuffle priorities before anyone clocks in. Geo-fenced timecards confirm boots are on site, and voice notes auto-transcribe into the master schedule—turning yesterday's static plan into a living system.

3-Step Live Workflow Setup

- Deploy a cloud platform that merges CPM schedules, cost codes, and mobile crew boards in one interface.

- Activate geo-fencing on time cards so payroll starts only when craft workers cross the gate.
- Turn on schedule recommendations; review daily “proposed shifts” instead of manually re-sequencing tasks.

Impact: Contractors unlock direct labor savings by compressing the wait-and-see gaps that traditionally consume 5–10 percent of field hours. Project managers and schedulers gain schedule transparency that they can forward to owners and public stakeholders.

Visual verification at production speed – Powered by Computer Vision

Progress reporting has long been hampered by subjective walk-throughs and after-the-fact punch lists. Computer vision engines now ingest drone imagery, mast-mounted cameras, and helmet-cam feeds, comparing as-built conditions to BIM in near real-time. The software flags out-of-tolerance welds, missing embeds, or misrouted conduit before downstream trades mobilize.

- Integrate drone and stationary camera feeds directly into your field data environment.
- Map photo IDs to model elements so that exceptions land on the exact task and cost code.
- Automate “ready for next trade” notifications once the dependent tasks are completed.

Pain point addressed: Rework, which typically accounts for 2–6 percent of project cost, is detected and corrected early, preserving both margin and schedule buffer. Warranty exposure drops because photographic evidence of compliance is updated.

Sensor-enabled data decisions replace guesswork

Embedded IoT sensors now cover concrete pours, MEP systems, and heavy equipment. Data from these can be integrated into mobile apps and analytics to provide feedback and foresight. From logistical production planning to curing timelines, sensors and tags can elevate site management.

Predictive Sensor Rollout Guide

- Use mobile apps to read sensors, RFIDs, and BLEs for onsite updates.
- Real-time cloud connections can immediately process these to give instant feedback and identify issues.
- A connected schedule can impact upcoming field tasks and the critical §§path.

Why it matters to executives: IoT-enabled projects can reduce costs by up to 29%. This significant cost reduction is achieved through enhanced real-time data exchange, remote monitoring, and improved project efficiency and safety. The accuracy of site data improves the quality of decisions.

Safety automation and compliance intelligence

Computer-vision models detect missing PPE, edge exposure, and unsafe proximity to operating equipment, triggering real-time alerts and logging incidents for trend analysis. Facial-recognition turnstiles verify worker credentials, ensuring site access safety. Impact sensors inside helmets feed data to AI models that distinguish genuine impacts from dropped gear, accelerating root-cause analysis without false alarms.

Smart Safety Deployment Steps

- Link AI safety cameras and smart PPE sensors to your live site map; violations surface as geotagged pins.
- Automate corrective-action workflows so foremen receive instant tasks instead of weekly safety reports.
- Sync credential databases with access control to stop out-of-date workers at the gate.

Risk-management payoff: Lower incident rates translate into favorable EMR scores and reduced insurance premiums, while owner confidence grows when AI-validated safety metrics are reported daily instead of monthly.

Jobsite connectivity that never blinks

Digitization is only as strong as the network spine that supports it. Private 5G, CBRS, and mesh Wi-Fi maintain high-bandwidth links from subterranean parking decks to rooftop penthouses, ensuring the AI models at the edge receive fresh data seconds after it’s generated. Offline-first apps buffer data when the signal drops, then sync automatically, keeping AI predictions intact and eliminating version conflicts.

Always-on network blueprint

- Stand up a site-wide mesh or private cellular network during mobilization, not after issues appear.

Metrics That Win Boardroom Approval

KPI	Typical Baseline	Digitized Field Target	Source of Improvement
Labor productivity (hours/CSI unit)	1.00	1.12–1.18	Optimized task sequencing & reduced waiting
Rework cost (% of contract)	3–6 %	<2 %	AI computer-vision detection & AR validation
Schedule variance (critical-path days)	10–15	<5	Change propagation
Lost-time incident rate (per 200 k hrs)	2.5	<1.5	Safety analytics
Equipment downtime (% of shift)	8–12 %	≤5 %	Predictive telematics & autonomous correction

- Use offline-first mobile apps that queue data locally and reconcile versions on reconnect.
- Monitor bandwidth health in the same dashboard as schedule and safety to catch dead zones early.

For the project office: Real-time analytics enable rolling-wave planning and same-day cost forecasting; disputes over "who knew what, when" lose traction because every stakeholder sees the same time stamped record.

Augmented field enablement

BIM viewers and augmented-reality overlays are no longer pilot curiosities. Field crews superimpose design intent over physical installations via tablets or smart glasses, while AI object-recognition confirms that fasteners, sleeve locations, and torques align with spec. Digital checklists surface inspection priorities, letting teams focus on the riskiest elements first. Completion evidence is logged with geotags and timestamps, feeding QA dashboards automatically.

AR/BIM Field Enablement Steps

- Push lightweight BIM models to field tablets with offline caching for sublevel work.
- Pair AR overlays with AI object recognition to verify component placement instantly.
- Auto-generate QC reports the moment crews tag "complete," reducing inspector wait time.

Operational upside for trade contractors: Quality becomes "built-in" rather than inspected in. Trade contractors reduce punch-list exposure, and GCs accelerate TCO hand-over because project data drops directly into the owner's AI-ready FM platform.

Bringing It All Together: The AI-Integrated Data Spine

True field digitization means one CDE plus AI orchestration: schedule data streams into safety dashboards, cost reports, and owner portals; AI cross-checks shifts against material call-offs, 3-D progress, earned-value curves, and safety heat maps, flagging clashes in real time.

Executives run a live control tower, not a weekly look-back, while AI classifiers turn every sensor feed into an ever-fresh digital twin, handing owners an operational model on day one.

Implementation Roadmap for the C-Suite

1. Audit the Data Landscape – Catalogue manual, semi-digital, and automated feeds; score each for AI readiness.

2. Prioritize AI Quick Wins – Mobile scheduling, computer-vision QA, and AI-sensor curing often pay back within one project.

3. Invest in Robust Connectivity – Private 5 G or CBRS is the oxygen your analytics need—budget for it early.

4. Standardize Data Schemas – Adopt ISO 19650 or similar frameworks so BIM, IoT, cost, and safety data interoperate for AI processing.

5. Pilot, Measure, Scale – Run controlled AI pilots, benchmark KPIs, and hard-gate rollouts on quantified gains.

6. Upskill the Workforce – Pair craft training with data & AI literacy; frontline adoption accelerates when crews see direct benefits.

Metrics That Win Boardroom Approval Position AI-enriched hand-over data as a value-add for facility managers, strengthening future bid competitiveness.

Field digitization has moved from pilot to board-level mandate, delivering hard gains in margin, risk control, and client trust. By unifying real-time data across materials, labor, safety, and equipment, contractors stay agile amid supply shocks, talent gaps, and stricter regulations.

At the same time, rivals drift into a credibility gap with owners who now demand data-proof. The frontier isn't another gadget; it's a single, intelligent backbone that orchestrates every hour, cubic yard, and safety check consistently finishing projects on time, on budget, and with unmatched transparency and resilience.

How the UK can unlock the power of renewables to fuel data centre potential

Data centres' energy needs are extensive and rising. They are also unique. To support them, and the business and growth they represent for the UK, we need to pioneer ways to leverage renewables to the maximum to support their intense energy needs.

BY TONY O'CARROLL, CEO, CONRAD ENERGY

ONE OPTION for data centres is to install behind the meter energy: power generated from on- or near-site sources, for instance wind turbines in a nearby field or solar panels on roofs. These can be installed under a Power Purchase Agreement (PPA), with energy companies covering upfront costs of installation and then selling the renewable energy back to the data centre at a fixed price for a fixed duration, whilst avoiding the non commodity charges associated with importing power from the grid. For those with more limited space, Corporate PPAs can be used instead to buy renewables from assets located elsewhere in the grid network. Both these options provide valuable benefits in terms of price and supply security, as well as providing sophisticated data for optimising and proving renewable usage.

Alternatively, another route is buying green power from the grid on standard supply contracts. However, in this respect fluctuating climate and infrastructure geared towards the patterns of fossil fuels are presenting challenges to harnessing sustainable energy with enough stability to be reliable.

What can be done to unlock this game-changing source of power from the grid?

Climate change putting heat on the grid

Encouragingly, production from renewable technologies increased to a record share of 50.8% of electricity generation last year. But the inherent intermittent and unpredictable nature of renewables is hindering the reliable

integration of these sustainable sources into the grid.

In the UK we are used to the challenges of Dunkelflaute – times of low renewable output in winter. On the flipside, June 2025 was the warmest on record – however this may not bring as many benefits for renewable usage as may instinctively seem to be the case. The collision of intense heat and a lack of wind (known as Hitzeflaute), and sudden spikes in both, has put the grid under stress, whilst a drop in sun in the evenings forces back traditional reliance on fossil fuels at these hours. The result? Fast-rising power prices and unstable power infrastructure.

With climate patterns shifting, it is now critical that the grid can adapt and flex to variable weather and power generation. If not, the energy transition will simply be unsustainable. To fuel the expansion of much-needed data centres, we cannot let this be the case.

What does our infrastructure currently look like?

The key to grid management is inertia – a force that works to prevent sudden frequency changes such as those that can be created by renewable generation. These sudden changes can cause blackouts like the ones recently suffered in Spain and Portugal. Data centres should not be operating with the threat of such events hanging over them – inertia needs to be in place.

To date, during the years of majority fossil fuel usage, the UK relied on synchronous generators at fossil fuel plants to provide this inertia and keep the grid balanced.

Alongside these, another important and historic component of our grid system has been the process of issuing grid connections. The unvetted first come first served system for this has created a logjam of zombie projects, unlikely ever to get off the ground, blocking the queue for viable – and vital – developments like the construction and connection of more data centres.

How can this ageing infrastructure be brought into the renewable world of the 21st century?

Science and strategy

In welcome news, the UK has been making solid progress in developing

infrastructure that can support the increasing integration of renewable power into data centre operations.

From a legislative perspective, the recent Labour Infrastructure and Industrial policies provide welcome promises of accelerated planning decisions and connection allocations for data centres.

Crucially, this is also married with activity on-the-ground in tech and innovation. Indeed, it's all very well boosting access to the grid, if the grid itself is not up to standards. To this end, we are working to replace the synchronous plants with synchronous condensers – rotating machines that provide inertia without generating power. These will help stabilise frequency fluctuations as traditional fossil plants are phased out. Grid-forming battery inverters are also a modern equivalent to historic synchronous machines, enabling batteries to provide system stability in real time.

It is this combination of science and strategy which is needed to power the grid forward and provide a proper foundation for the data centre industry.

Renewing our investment in renewables

There is, of course, still more that needs to be done. We need to tackle grid infrastructure at every level – including improving local grids. Here, boosting our battery infrastructure will be essential – Lithium-ion batteries are compact enough to be transportable and installable in locations to ensure flexibility at local levels.

At the other end of the spectrum, we also need to scale these possibilities up to bolster not only flexibility, but also energy storage, on a national level. Pumped hydro and green hydrogen are options, but again it is Lithium-ion batteries which currently provide the most timely, expandable and financially feasible solution in the foreseeable future.

If a robust contractual framework, reinforced by compelling incentives, is put in place to help develop battery infrastructure, it would make more commercial sense to build sites that store renewable energy for far longer than the 1-2 hours which is currently possible. This will maximise the efficiency with which we use our renewably generated energy - allowing us to retain excess power and release it back into the grid when generation drops but demand rises.

For data centres, these steps will be crucial to support the stability of the grid and in turn access to renewable energy. The industry will be waiting with bated breath for this progress to be made.

Sustaining data centre growth

The widespread roll out of renewable energy was once seen as a far-off dream, but the fact is this is fast becoming reality. Data centre operators should feel empowered to seek out renewable solutions for their energy usage – but it is essential we have wider infrastructure that can keep up with these ambitions.

We are already seeing this integration take place, with promising results. It is essential the UK now builds on this foundation to remove the stability, predictability, and storage barriers which remain, and allow these sustainable sources to power the digital heart of the UK: our data centres.

Pros and cons of Small Nuclear Reactors in the datacentre construction industry

Following a competitive two-year bidding process with two US rivals, Rolls-Royce was recently awarded the contract by state-owned Great British Energy - Nuclear: to become the first company to build small modular nuclear reactors (SMRs) in the UK.

BY SEBASTIAN MURPHY, TECHNICAL DIRECTOR - DATA CENTERS EMEA, BLU-3, SHAHEED SALIE, TECHNICAL MANAGER, BLU-3, AND WARREN MCTACKETT, BLU-3.

ALONGSIDE £14.2bn of investment already pledged for the Sizewell C power station in Suffolk, the advent of domestic SMRs, to which the government is pledging over £2.5 billion, is part of its wider effort to position Britain at the forefront of nuclear energy technology. It is anticipated that Sizewell C will produce 3.2 GW of power, while each of the three Rolls-Royce's SMRs will provide about 470 MW.

SMRs are seen as a potential power source for datacentres because they provide a reliable, low-carbon, and scalable energy supply. Globally, the exponential growth in demand for AI-related services is driving demand for AI data centres that require consistent power delivered by SMRs.

Companies such as Google, Amazon and Oracle have recently announced plans to use SMR energy to run their data centres. To manage the huge technological demands, AI needs bigger, more complex and power-hungry data centres, which supply chains and national power grids may find challenging to meet.

The news about Rolls-Royce building SMRs in the UK is welcome. Based on our project teams' on-site experience across EU member states, and our preconstruction activities, such as viability, feasibility, and buildability analysis, it does raise several challenges in relation to how they will work in practice to service datacentres.

Critically, the SMR model is yet to be commercially demonstrated: no sites are yet fully operational anywhere in the world.

The first point of concern is the temporal disconnect that exists in project timelines. The construction and commissioning of nuclear facilities involve lengthy processes - typically, ten to 15 years - whereas the construction of hyperscale facilities operates on a much shorter timeline. In an ambitious statement, the UK government hopes 'to connect the SMR projects to the grid in the mid-2030s.'

Hyperscale facilities are designed to handle enormous amounts of data processing, storage, and computing power. In our experience, they usually have a three-to-five-year planning to completion cycle. By the time the first SMR is scheduled for completion in the UK, the data centre's tech stack may be more than a decade old, making it potentially less aligned with future requirements.

The next point to consider is data centre consumption, which is based on dynamically fluctuating loads. This means it can transform from being idle to spiking upwards at any point: they are designed to scale up or down rapidly in order to meet fluctuating demand.

By contrast, the SMRs provide a consistent and constant output. Because their capacity to supply these types of facilities has not yet been demonstrated, it is therefore advisable that the load capabilities of SMRs are investigated during the early stages of design. Notably, nuclear energy may complement existing backup systems, such as diesel gensets, UPS systems, and battery backups.

Cost and return on investment (ROI) are always critical challenges in the evaluation of any large-scale infrastructure project: determined by the profitability of an investment proportionate to the total costs involved.

The construction of data centres is predicated on a rapid ROI model, whereas nuclear energy is based on long-term, high capital expenditure. Without long-term power purchase agreements (PPAs) or state-backed incentives, it is challenging to synchronise the respective investment cycles of nuclear developers and hyperscale clients because they are so different.

Public perception remains an ongoing topic of consideration in the nuclear industry. History demonstrates that nuclear power has generated public scepticism over time: fuelled by concerns over safety, due to potential risks during operation and the management of the waste, and potential links to nuclear weapons proliferation.

Public perception remains an ongoing topic of consideration in the nuclear industry. History demonstrates that nuclear power has generated public scepticism over time: fuelled by concerns over safety, due to potential risks during operation and the management of the waste, and potential links to nuclear weapons proliferation

Manifestly, there will be challenges in maintaining public confidence when adopting nuclear power as a source of energy for datacentres, which can often be contentious within the communities where they are built.

For example, public demonstrations have recently occurred in the Netherlands, involving local farmers and environmental groups, over proposed government net zero policies aimed at reducing nitrogen emissions and expanding nuclear energy.

Location is another key issue. Any decision to locate data centres adjacent to small nuclear reactors introduces important considerations and, since the UK lacks clear rules for direct (behind-the-meter) connections, several outstanding details to be clarified: who supplies the power, how are costs set, and who holds liability?

Having sufficient space to guarantee safety is a crucial consideration. SMRs need to create safety zones, cooling systems, and security buffers: the space that is required to fulfil these objectives makes it more complex to integrate close to dense urban or fibre-rich sites where data centres usually operate, such as edge data centres.
As outlined above, waste management is a central focus in relation to nuclear projects.

All reactors, including the planned SMRs, produce nuclear waste. Long-term storage and handling (even if it is reduced) will require national policy alignment and appropriate facility agreements to be implemented.

Elsewhere, alternatives to nuclear power for use on datacentres are being investigated. The Netherlands is looking at the potential adoption of hydrogen energy, because its combustion produces only water, without any direct carbon emissions.

As a key part of its transition to sustainable energy, the Netherlands is investing heavily in hydrogen energy, particularly green hydrogen. This would enable excess power from local wind turbines to be used in order to convert water to its component elements - hydrogen and oxygen. Because the Netherlands has an abundance of both wind and water, it could become a major competitor in the field of alternative energy sources, particularly for datacentres.

Although there are logistical and technical considerations for implementing these green benefits, the use of hydrogen energy clearly has significant potential to reduce greenhouse gas emissions.

The Dutch government's vision also aligns with many of our tech clients' goals to become entirely carbon neutral, or even negative, within the next decade or two. Equally, any scientific advances in energy output that could potentially mitigate nitrogen oxides emissions on datacentres will also help the industry to comply with EU legislation and regulations.

Although it requires thoughtful planning relating to waste management, and the substantial cost and long-term planning required for construction, nuclear power is also a low-carbon energy source: it does not produce greenhouse gas emissions during electricity generation.

Even though key advantage of nuclear is that it only requires a small amount of fuel to provide power for a significant duration: as a fuel source, it could potentially outlive the lifespan of server rack inside a datacentre to which it provides energy.

SMRs will be much smaller than conventional nuclear reactors. Potentially, they can be built in factories. As a faster and more cost-effective way to deploy nuclear power, their modular, factory-based construction is a key advantage. Occupying a smaller footprint than traditional nuclear plants, they can be modularly scaled, making them more adaptable to phased datacentre expansion plans. Co-located SMRs will also reduce reliance on national grids, improving resilience against outages, grid congestion, or power price volatility.

Although cheaper and quicker to deploy than large nuclear power plants, multiple SMR designs exist. 68 are currently listed by the International Atomic Energy Agency (IAEA). Ultimately, their individual success will depend upon demonstrating that they are fit for purpose.

Data centre location strategy in the Age of AI

Investment headlines still prominently feature traditional data centre hubs such as Northern Virginia, Frankfurt, and Singapore. Yet beneath these familiar names, data centre developers are navigating a rapidly evolving landscape driven by unprecedented shifts in technology, capital, and regional constraints.

BY AASHNA PURI, GLOBAL STRATEGY AND EXPANSION DIRECTOR, CYRUSONE

THE EXPLOSIVE growth of artificial intelligence (AI), combined with enormous capital requirements and complex physical limitations of power and land, is redefining how and where data centres are built, ushering in a dynamic era of growth where uncertainty is not just a challenge, but a fundamental part of strategic decision-making. Location strategy now relies on three forces working together: the workload in scope, the capital required to deliver at speed, and the regional constraints that ultimately decide what is possible.

AI and the shift from predictability to rapid scale

Traditional data centre development relied on steady, predictable growth of enterprise and cloud workloads. Facilities could be planned with relative confidence, expanding incrementally to match demand. That model alone though no longer works.

While enterprise and cloud adoption continue to grow, AI has added a new layer of location criteria driven by the very different needs of training and inference workloads.

Training, used for building large language models and generative AI, requires massive, centralised campuses with extremely high-density compute. These facilities are power-intensive, thermally complex, and increasingly shaped by rapid hardware innovation.

AWS, for example, recently reported that its custom Trainium2 chips now offer 30-40% better price-performance than GPUs. Shifts like this can alter power, cooling, and rack density requirements mid-cycle. Because training is far less sensitive to latency,

it gravitates towards lower-cost, power abundant markets where hundreds of megawatts can be assembled at a low total cost of ownership.

On the other hand, inference workloads support real-time user interactions and require low-latency delivery close to end users. Here location strategy revolves around major population centres with strong connectivity through dense fibre networks, cloud on-ramps, and local internet exchanges, while still balancing power availability. Regulations and data residency requirements can also dictate where these sites need to be.

In practice, developers now serve three distinct, sometimes overlapping demands: enterprise workloads in availability zones, large, centralised campuses for training, and distributed infrastructure for inference. Building quickly remains essential, but success depends on being able to adjust capacity and design as workloads change.

Capital: fuelling the scale of AI expansion

To match the scale of AI-driven growth, unprecedented levels of investment are required. Morgan Stanley forecasts global data centre capital expenditures could reach nearly $3 trillion by 2028, a number far exceeding any previous infrastructure cycle, including the late-90s telecom boom. While hyperscalers like AWS, Google, and Microsoft, will directly fund much of this, private credit will be critical in bridging an estimated $800 billion gap. Debt markets are also expected to triple the size of today's ABS and CMBS market through $150 billion in bonds backed by data centres and related hardware.

However, this capital comes with expectations. GPUs can lose as much as 30 percent of their value each year, and with rapid hardware turnover, shifting power needs and evolving design standards, investors will favour facilities that can be adapted, upgraded and reconfigured as technology changes. Projects with ready power, land and permits are best placed to deliver returns quickly.

As a result, capital is not only enabling the next wave of growth but also influencing which locations rise to the top of the build queue.

Hyperscalers are building gigawatt-scale campuses to support AI training at unprecedented scale, while model builders, neocloud platforms, and chipmakers are all competing for power and land infrastructure that can meet their specialised needs

Regional realities: balancing land, power, and regulation

Beyond technology and capital, there are physical constraints like power availability, and evolving regulatory frameworks that impact where data centres can be built. Each region has distinct strategic challenges and opportunities.

In the United States, where more than half of all new global capacity is expected, demand is being driven by a broad mix of customers and workloads. Hyperscalers are building gigawatt-scale campuses to support AI training at unprecedented scale, while model builders, neocloud platforms, and chipmakers are all competing for power and land infrastructure that can meet their specialised needs. While established markets like Northern Virginia face space and power constraints, this is accelerating growth in secondary markets like Reno, Columbus, and the broader Dallas–Fort Worth region. These locations offer the land, power, and permitting speed needed for rapid, large-scale deployment – reinforcing the US as the centre of gravity for global AI infrastructure.

In Europe, the data centre landscape is evolving through strategic layering of workloads. Growth continues steadily in enterprise cloud adoption, with AI workloads increasingly being integrated into existing facilities. This convergence optimises existing infrastructure, leveraging colocation to balance enterprise and AI demands within the same locations. Demand is also being shaped by regulatory compliance and sovereignty requirements, particularly in sectors like government, finance, and healthcare, where data residency and low latency are critical. As a result, primary markets continue to see the bulk of activity, positioning Europe as a leader in sustainable data centre development.

Asia, contrastingly, has a hub-and-spoke model: core hubs such as Singapore, Tokyo, and Seoul handle critical workloads, while constrained land and power push new capacity to nearby spokes like Johor, Batam, Osaka and Busan. Further, fragmented sovereignty rules are forcing operators to replicate smaller footprints in each market, resulting in networks of medium scale, low-latency sites over single mega-campuses. Overall, strong local policy and capital support help projects move comparatively quickly, making Asia's market more nationally segmented than Europe and far more geographically dispersed than the gigawatt-scale developments seen in the US.

Future-Proofing Location Strategy: Balancing Speed, Capital, and Change
Ultimately, the interplay between rapidly evolving AI technologies, extraordinary capital requirements, and intricate regional constraints calls for a holistic and adaptable approach. Data centre strategies must evolve from predictable, incremental growth models toward flexible frameworks designed to rapidly respond to technological advancements and shifting market conditions.

In the AI-driven future, success for data centre developers will hinge not merely on anticipating growth but on effectively integrating these interconnected dynamics. It will mean building infrastructure capable of meeting demand for centralised, energy-dense requirements of AI training and distributed, latency-sensitivity needs of inference workloads.

It will mean selecting strategic locations where capital, regulation, land availability, and power align favourably. Most importantly, it will require embracing uncertainty as an inherent part of strategic decision-making.

Enhancing hyperscale strategies with edge data centres

Digital infrastructure is undergoing a quiet revolution. As demand for real-time responsiveness, massive data throughput, and local compliance grows, the traditional hyperscale model - relying on a few centralised mega-data centers - is reaching its limits. Meanwhile, the edge is emerging not as a competitor, but as a complementary force. When combined, edge data centers, together with hyperscalers can deliver the flexibility, reach, and performance required by modern digital services.

BY HANS NIPSHAGEN, VICE PRESIDENT CHANNEL AT NLIGHTEN

OVER the past two decades, hyperscalers have built their success on highly centralised infrastructure - large data centers positioned in strategic network and power hubs. This model has facilitated rapid global scale and operational efficiency. Yet as digital services evolve, new pressures are emerging.

Applications that depend on ultra-low latency, high availability, or local data processing increasingly expose the limitations of centralised architectures. With the mainstream adoption of smart manufacturing, autonomous systems, and AR/VR, even milliseconds of delay can undermine functionality.

Additionally, constant backhauling of vast amounts of data to centralised sites strains networks and unnecessarily increases energy consumption. Data sovereignty regulations further compel hyperscalers to adopt more localised approaches.

Edge computing is bringing digital infrastructure closer

Edge data centers – typically smaller, distributed facilities deployed closer to users and devices – help alleviate these challenges. By reducing the distance data must travel, they reduce latency and enhance application responsiveness.

This is particularly beneficial for real-time use cases such as predictive maintenance, intelligent transportation, or high-frequency trading. Beyond performance enhancement, edge computing serves as a strategic enabler. In regions with limited network infrastructure, localised computing bridges service gaps and ensures consistent digital experiences. This empowers hyperscalers to tap into new markets and user segments that would otherwise remain out of reach.

The synergy: Hybrid infrastructure in action

The most effective strategy isn't to replace one model with another, but to integrate them. Hyperscalers can retain their centralised core for heavy-lift workloads and archival storage, while delegating latency-sensitive and localised tasks to the edge.

This hybrid architecture supports dynamic workload placement, stronger disaster recovery, and improved compliance across jurisdictions.

Sustainability and visibility

Sustainability is no longer a bonus but a baseline expectation, and here edge can make a meaningful difference for hyperscalers too. With the growing scrutiny on the climate impact of digital services, infrastructure operators are under pressure to reduce carbon footprints and improve energy transparency. Edge facilities, when thoughtfully designed, can help achieve both. Many are built with local power grids in mind, incorporating renewables and energy-efficient cooling solutions. Importantly, newer energy agreements provide operators with detailed, asset-level data about their power usage, offering unmatched internal visibility while supporting external requirements.

Looking ahead

Edge computing is more than a technical upgrade; it is a strategic lever for hyperscalers to enhance their core offerings, stay ahead of regulatory trends, and unlock new value. Building hybrid infrastructure today positions hyperscalers to meet the demands of tomorrow, from AI at the edge to local digital ecosystems. The future of cloud infrastructure lies in embracing both scale and proximity. By complementing their centralised platforms with edge deployments, hyperscalers can deliver faster, smarter, and more sustainable services anywhere users need them. This is not about shifting paradigms. It's about building on the existing one.

TRILUX E-LINE PRO

THE LIGHTING SOLUTION FOR DATA CENTRES

DATA CENTRES ARE GROWING RAPIDLY – AND WITH THE E-LINE PRO, TRILUX IS DELIVERING THE IDEAL LIGHTING SOLUTION FOR DATA CENTRE HALLS.

High-performance optics and state-of-the-art LED technology ensure homogeneous light with excellent colour rendering (CRI >90) on vertical surfaces. This increases safety and efficiency when working. Our ergonomic systems score with high energy efficiency, low heat input and a long service life. Also integrable: emergency lighting, data cables, sensors.

Intelligent light management can be included – for minimal maintenance.

www.trilux.com/data-center

The shift from standby to strategic energy management

It's safe to say the energy landscape is changing. When I think about my 5-year career working with data centers, the most prominent and significant changes have taken place in the last 24 months. The data-driven society we live in, from streaming devices and smart appliances to AI processing, continues to move demand for data centers in just one direction: up.

BY LAURA MACIOSEK, DIRECTOR KEY ACCOUNTS, CAT® ELECTRIC POWER DIVISION

IT'S SAFE TO SAY the energy landscape is changing. When I think about my 5-year career working with data centers, the most prominent and significant changes have taken place in the last 24 months.

The data-driven society we live in, from streaming devices and smart appliances to AI processing, continues to move demand for data centers in just one direction: up.

As data centers experience this growth, utility power is no longer a given. Today, there's no guarantee the local electrical grid can meet these increased power needs. In fact, many utilities I've talked with say it'll be three to five years (or longer) before they can bring the required amount of power online. That puts data center customers in a tricky position. How can they continue to expand and grow if there isn't enough power and moving sites isn't an option?

The answer includes rethinking power options, and that means considering the transition from using power assets for largely backup purposes to employing them as a primary power source. That's a big change from the status quo. If you're in a similar position, here's some advice on how to navigate the transition.

Bridge The Gap

First things first: if you need power quickly, and your utility cannot provide it, you should be considering a "bridge to grid" solution. This refers to "bridging the gap" until the local grid can provide sufficient power. These solutions can be made up of multiple power generation assets — including diesel and natural gas generator sets or turbines.

Why this works

The beauty of a bridge-to-grid option is that it takes only a matter of weeks to install and can remain in place for weeks, months or years. That allows you to set up or expand a data center operation, address growing power requirements and ensure reliability in places you wouldn't have been able to do so before.

Consider your emissions requirements

One concern I hear from customers embarking on this change is the impact on emissions. Switching from standby to prime power requires you to run your power generation assets more frequently, and that can affect your site's overall greenhouse gas (GHG) output — a challenge for data centers tasked with meeting evolving 2030, 2040 and 2050 sustainability goals.

Good News: You Have Options

To support your sustainability journey and reduce GHG emissions, you can:

- Add (or retrofit) a Clean Emissions Module (CEM) to your diesel generator set to support Tier 4 emissions compliance, including reduced particulate matter and NOx emissions.
- Run your generator sets on biofuels, 25% hydrogen gas blends, or even 100% hydrogen.
- Install microgrid solutions where possible, maximizing your use of renewable energy sources.

Make the economics work

Right now, your focus is likely capacity, but keep in mind that a power solution can be more than just a resiliency play. If you plan properly, it can benefit you economically, too.

For example, once your site is connected to the grid, what is your plan?

One option is to transition your bridge-to-grid assets to a more traditional standby role and then install a distributed energy management system (DERMS) on them, which can save or even make your business money.

Why this works

With a DERMS in place, your power generation assets will be automatically dispatched when it makes the most financial sense, like during peak electricity usage when grid prices are at their highest.

This helps your business avoid peak charges and lower overall energy costs. There's also the possibility to sell excess energy back to the grid. Either option can help you offset the initial capital cost of equipment and generate long-term savings.

Create your long-term plan

Often customers ask me: How can I address my immediate needs without ending up in a similar situation in a few years' time? My answer: Plan 15 years out.

I encourage and work with customers to create a plan that considers:

- **Site requirements:** Do you have space to expand? Do you have space for renewable sources?
- **Growth projection:** How much do you need to expand in the next 15 years? How much energy will you need?
- **Emissions reporting:** What are your emissions targets? How can you achieve them?
- **Overall landscape:** What local requirements might you need to comply with in the future? Do you have a plan if that happens?

It is a time-intensive process, but once completed you'll have peace of mind that you should have enough power to see you through.

Whether you're ready to make the switch from standby to prime power at your data center today, or simply weighing options for your next development or expansion, we're here to help.

We'll work with you to find the right combination of assets and asset management software that fulfills your power requirements reliably and cost-effectively.

Connect with one of our experts to get the process started.

With a DERMS in place, your power generation assets will be automatically dispatched when it makes the most financial sense, like during peak electricity usage when grid prices are at their highest. This helps your business avoid peak charges and lower overall energy costs. There's also the possibility to sell excess energy back to the grid. Either option can help you offset the initial capital cost of equipment and generate long-term savings

Hybrid cooling explained: why one term doesn't fit all

As AI accelerates, high-performance computing becomes the norm, and sustainability moves from ambition to imperative, meaning data centres are under pressure to evolve quickly, intelligently, and responsibly.

BY STUART FARMER, SALES DIRECTOR AT MERCURY POWER

THERE'S a problem. Traditional cooling methods are nearing their limits.
In response, many operators are turning to hybrid cooling. It's a phrase that's gaining traction across the industry, but what does it really mean?

Ask five industry experts, and you'll get five different answers. That's because hybrid cooling isn't a product or a plug-and-play solution. It's a strategy, and if we keep treating it like a single technology, we risk missing its real potential.

What is hybrid cooling - really?

Hybrid cooling isn't just one system. It's a tailored blend of multiple cooling technologies, adapted to the specific needs of a site, its location, workload density, and sustainability goals. It's about building resilience and efficiency by design, not by default.
Let's take a look at what this looks like in the real world.

Air and liquid cooling – most common definition

Air cooling through chilled air or free cooling handles standard server loads. High-density racks (often AI-driven) are cooled using liquid systems, whether direct-to-chip or immersion. This setup is common among hyperscalers like Meta, Google, and Microsoft.

The hybrid element comes from the selective deployment of liquid cooling, targeting it where it adds the most value rather than using it universally.

Mechanical and free cooling systems

Mechanical cooling (like CRAC units or chillers) activates during warmer conditions, while free cooling, using outside air, takes over when the weather allows. It's widespread among traditional colocation providers.

The hybrid approach helps reduce operational costs and environmental impact by minimising reliance on refrigerants and compressors.

Evaporative and mechanical cooling

Evaporative systems use moisture to cool air, making them ideal for dry climates. Mechanical systems offer backup during periods of high humidity or when redundancy is needed. You'll often find this setup in hot, arid regions like the southwestern U.S., the Middle East, and Australia.

Water-based cooling with air distribution

Here, liquid systems draw heat directly from servers, while air systems manage any residual heat and maintain environmental balance. This approach is particularly popular in advanced European facilities - IBM can be seen using this method. By combining point-of-use efficiency with traditional room-wide control, this hybrid model strikes a powerful balance between precision and scale.

Onsite and district or external cooling

In this model, internal systems cool the data centre, and recovered heat is reused through district heating networks. An approach gaining traction in colder urban environments, such as Scandinavia and Northern Europe.

By feeding waste heat into city-wide systems, these setups reduce environmental impact and support circular economy models.

AI-driven dynamic hybrid cooling

Using machine learning, these systems anticipate thermal loads and environmental shifts, dynamically switching between modes for optimal efficiency.

You'll often find this approach in cutting-edge data centres with sophisticated energy and thermal management platforms. This isn't just hybrid in design, it's hybrid in motion. Intelligent, adaptive, and deeply responsive.

The downside of oversimplification

Hybrid cooling is not a checkbox. It's not something you "install." It's a mindset, a design philosophy that respects context.

When hybrid systems are poorly integrated, competing subsystems can cancel out each other's benefits. Overengineering can also inflate complexity and cost, undermining the original goals of resilience and efficiency.

The smart question isn't "What's the best hybrid?" It's "What's the right combination for us?"

Final thought

We believe in clarity over jargon and the industry needs this more than ever. That's why we cut through the noise to deliver cooling strategies that work - not just on paper, but in practice.

Hybrid cooling isn't the future. It's already here. But how we define it, and more importantly, how we design with it, will determine whether it simply keeps up with change or helps us lead through it. And it's only as good as the thinking behind it.

Data centre availability crisis deepens as vacancy hits historic low

JLL's new North America Data Center Report – Midyear 2025 reveals how "power is the new real estate" as the industry races to add capacity.

THE North America Data Center colocation market has reached a critical tipping point, with vacancy rates plummeting to an unprecedented 2.3% amid relentless demand for digital infrastructure. JLL's North America Data Center Report – Midyear 2025 reveals how, as inventory grows to a record 15.5 GW, the sector continues its explosive growth trajectory while grappling with severe capacity constraints and energy sourcing challenges.

Northern Virginia maintains its position as North America's the largest data center market. With 5.6 GW of capacity, it is more than triple the size of Dallas-Fort Worth, the second-largest market at 1.5 GW.
The report also highlights how cloud providers and technology companies continue to dominate data center demand, accounting for 65% of all leasing activity.

"The colocation market is experiencing unprecedented demand pressure under an increasingly stressful environment," said Andy Cvengros, Executive Managing Director, Co-Lead of U.S. Data Center Markets, JLL. "The first half of the year was riddled with disruptions, including the DeepSeek news at the beginning of the year and the potential impact of tariffs on demand and construction. Despite the turbulence, the sector posted another record-shattering performance."

Absorption outpacing supply in record market run

The market absorbed an impressive 2.2 GW in H1, with half of this activity concentrated in Northern Virginia (647 MW) and Dallas-Fort Worth (575 MW). Chicago (368 MW) and Austin/San Antonio (291 MW) also showed significant leasing activity, putting the sector on pace to exceed 2024's record absorption levels.

"What we're seeing across primary markets is nothing short of extraordinary," said Curt Holcomb, Managing Director with JLL's global Data Center Solutions practice team. "In Dallas-Fort Worth, for example, there is unparalleled competition for limited capacity. Major cloud providers are securing power reservations years in advance, and the development pipeline has expanded to over 1 GW under construction. Meanwhile, Austin has emerged as a genuine Tier 1 market with nearly 921 MW of inventory and another 341 MW under construction, representing a 500% growth since 2020."

The construction pipeline has ballooned to 7.8 GW, approximately 10 times the volume from five years ago. Phoenix (1.3 GW) leads development activity followed by Chicago (1.18 GW) and Atlanta (1.11 GW) leading development activity outside Northern Virginia.

More concerning for those seeking space is that 73% of all capacity under construction is already preleased. High preleasing has remained consistent for the past two years, signaling meaningful market relief remains years away.

"The days of build-it-and-they-will-come are long gone," said Matt Landek, Division President, U.S. Data Center Work Dynamics, who also leads JLL's Data Center Project Development and Services. "What we're seeing now is

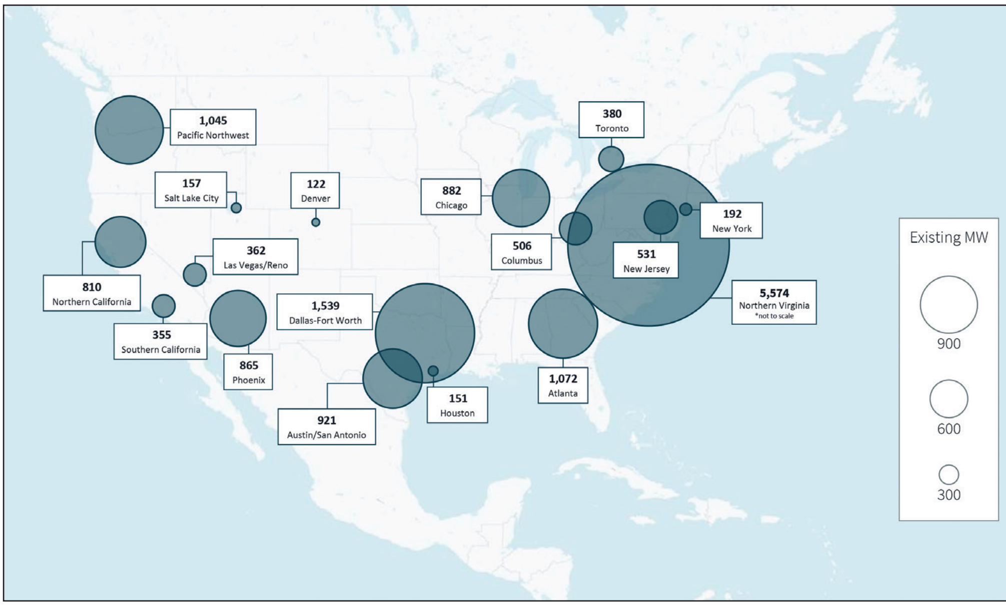

Midyear 2025 North America colocation statistics	
Inventory	15.5 GW
Vacancy	2.3%
H1 2025 absorption	2.2 GW
H1 2025 completions	1.6 GW
Under construction	7.8 GW
U/C % preleased	73%
Planned pipeline	31.6 GW
Average rents	($/kW/mo)
<250 kW (all in)	$327
250 kW-1 MW	$194
1-5 MW	$157
5-20 MW	$144
>20 MW	$129

North America colocation markets.

'commit-before-it's-built-or-you-won't-get-in.' This is fundamentally changing how companies approach their data center strategies. Enterprise users who once planned 6 to 12 months in advance are now securing capacity and their facilities and operations teams 18 to 24 months before their intended deployment dates, sometimes even earlier."

The high-voltage hunt for affordable power

While established markets continue to dominate the landscape, emerging markets are experiencing dramatic growth. Columbus has seen an astounding 1,800% growth since 2020, while Austin/San Antonio has grown 500% from a smaller base during the same period. These emerging markets are benefiting as power constraints in primary markets push development elsewhere.

Commercial electricity rates have risen nearly 30% since 2020, reaching an average of 9.7 cents/kWh in H1 2025. This increasing cost pressure has driven development toward markets with lower power costs such as Salt Lake City (5.7 cents/kWh) and Denver (6.4 cents/kWh).

The average wait time for a grid connection across the U.S. is now four years. Power delays remain a significant hurdle in alleviating supply constraints. However, there is a silver lining, this obstacle is also preventing a bubble from forming in the sector.

"Power has become the new real estate," said Andrew Batson, Head of U.S. Data Center Research at JLL. "With vacancy effectively at 0%, virtually all absorption is the result of preleasing with delivery times extending beyond 12 months. The market has been growing at a remarkable 20% CAGR since 2017, and our development pipeline data suggests this pace will continue through 2030, with the colocation market potentially expanding to 42 GW of capacity."

Investors double down on data centers

The data center sector continues to solidify its position as one of the most favored real estate asset classes, with market capitalization growth of 161% since 2019, second only to industrial properties. This remarkable growth is driven by insatiable tenant demand, limited supply and rising rents, creating a compelling investment thesis that continues to attract new capital to the sector.

The debt markets for data centers are experiencing significant expansion, with asset-backed security (ABS) and single-asset single-borrower (SASB) loan activity increasing for the third consecutive year. The first half of 2025 saw 14 ABS deals totaling $7.7 billion and four SASB deals totaling $5.7 billion, substantial increases from the same period in 2024. Meanwhile, asset-level investment sales remained relatively muted at $754 million across 23 transactions in H1 2025, with average cap rates holding steady around 6%, comparable to premium industrial and multi-housing properties.

Charting the course to 2030

Looking ahead, JLL anticipates the supply-demand imbalance to persist over the next several years. Projects already under construction are 73% preleased, and while an additional 31.6 GW of capacity is planned, that supply will be phased over five years or more. Northern Virginia leads all markets with 5.9 GW planned, followed by Phoenix (4.2 GW), Dallas-Fort Worth (3.9 GW) and Las Vegas/Reno (3.5 GW).

"North America could see $1 trillion of data center development between 2025 and 2030," Batson added. "Based on our forecast, more than 100 GW of colocation and hyperscale capacity could break ground or deliver over the next five years. These projections do not include the potential upside of quantum computing, which we see as a sector accelerant over the next 5 to 10 years."

The combination of AI adoption, digital transformation initiatives and cloud migration has created a perfect storm of demand that the industry simply cannot meet quickly enough, leading to the current supply crunch and making forward planning more crucial than ever for enterprises seeking data center capacity.

DENALI™ Optical Fiber Platform: How modular design meets the rising demands of AI and hyperscale networks

The growing global appetite for ultra-fast, high-capacity fiber networks capable of managing complex AI workloads continues to drive innovation in data center hardware design. With the increasing adoption of GPU-intensive AI clusters, hyperscale cloud environments must now embrace scalable, flexible fiber deployment options, suited to future growth. In response to these evolving requirements, AFL has introduced the DENALI Optical Fiber Platform, engineered to accelerate deployment times, simplify infrastructure oversight, and elevate fiber network performance.

Why DENALI Aligns with AI-Focused Network Growth

To maintain peak performance in fiber network environments, the physical layer must evolve to support increased density, broader bandwidth, and lower latency. Innovation in data center hardware and infrastructure planning must also facilitate rapid, streamlined, single-person integrations, ensuring minimal disruption and faster go-live times.

While traditional data center deployments could fail to adequately accommodate these ramping expectations, DENALI was specifically engineered to meet tomorrow's data center challenges head on – the platform not only enables seamless, modular expansion but assists installers by optimizing cable management in space-limited environments.

How a modular design enables scalable, high-density fiber infrastructure

DENALI features a modular, tray-based system compatible with fanout, patch, and splice cassette formats – all within a unified housing footprint. Supporting up to 288 LC Duplex Ports (576 fibers) in just 4RU, with 2RU and 1RU options also available, the DENALI platform delivers high-density fiber management tailored for modern data centers.

Key features include:

- **Universal Tray**
  One tray design supports all cassette types, simplifying ordering, installation, and future upgrades.
- **Integrated Cable Management**
  Rear trunk routing promotes clean, performance-optimized installations.
- **Detachable Front Patch Cord Clips**
  Improves accessibility and simplifies maintenance during service operations.
- **Front and Rear Cassette Access**
  Locking sliding trays allow flexible access, ideal for tight spaces.
- **Magnetic Front Panel**
  Provides secure access without hinges or latches.

These design elements support faster installations, lower labor costs, and simplified scaling. The DENALI optical fiber platform was built with rapid deployment and strong infrastructure ROI in mind.

Supporting Latency-Sensitive, GPU-Driven AI Workloads

AI workloads continue to push bandwidth requirements to unprecedented levels. Each new generation of high-speed GPUs and interconnect technologies raises the bar for infrastructure performance, positioning fiber networks as strategic assets in AI and cloud environments. DENALI was engineered to meet these demands, supporting bandwidths from 10G to 800G and beyond. This capability supports efficient AI cluster scaling, while a modular approach shortens deployment timelines and simplifies expansion, maintenance, and upgrades.

"The market is undergoing a major shift, where AI-driven densification is transforming how data centers approach fiber deployment. The DENALI platform was developed in response to this shift of handling faster scaling, reduced downtime and solid reliability that AI workloads actually need."

Marc Bolick, President of Product Solutions, AFL

Global Deployment Backed by Regional Compliance Expertise

AFL's global reach accelerates delivery times and provides end-to-end logistical support throughout every phase of deployment. DENALI empowers network operators to implement a standardized fiber platform worldwide, while ensuring compliance with regional standards.

This blend of global scalability and localized execution makes the platform the ideal choice for distributed architectures and geographically diverse data centers.

Plug-and-Play Fiber Ecosystem for AI Infrastructure

DENALI's single-person installation model extends across the entire product line, offering dependable plug-and-play solutions for modern fiber ecosystems. A comprehensive accessory portfolio supports rapid Moves, Adds, and Changes (MACs) and accelerates time-to-revenue for new builds.

The DENALI Ecosystem Includes:

- **Fiber Housings**
  Core platform component supporting all cassette types in a universal tray system. Up to 576 fibers in 4RU, with 2RU and 1RU options available.
- **MPO Fanout Cassettes**
  High-performance plug-and-play units supporting up to 24 fibers per cassette in Base-8 to Base-24 formats. Features include shuttered adapters and internal polarity management.
- **Patch Cassettes**
  Enable efficient cross-connections in dense environments, compatible with all DENALI housings.
- **Splice Cassettes**
  Designed for clean, efficient splicing with pre-routed pigtails for simplified field termination.
- **Outback Clip Management (OCM) Bracket**
  Organizes and secures high-fiber-count trunk cables, supporting up to 12 Outback Clip-mounted trunks.
- **MPO Trunk Assemblies**
  Built with AFL's MicroCore® reduced-diameter cable for enhanced bend tolerance and airflow.
- **LC Patch Cord Assemblies**
  Compact, high-performance cords with field-reversible Uniboot LC connectors for dense deployments.
- **Accessory Range**
  Includes mounting brackets and conversion kits for flexible deployment across evolving environments.

Strategic Infrastructure for Scalable Data Center Operations

Large-scale data center success depends on forward-thinking infrastructure choices. By consolidating part numbers, DENALI simplifies procurement and inventory management. Single-person installation reduces labor needs, while the accessory suite ensures seamless integration across diverse network environments. DENALI provides a flexible foundation to support every stage of ongoing network evolution. Stay informed at AFLglobal.com, follow us on LinkedIn, or Contact Us.

How AI is changing Data Centre power, cabling, and connectivity requirements

AI isn't just changing what runs inside data centres—it's reshaping the buildings themselves. Training and serving AI models are driving a step-change in electricity demand. Credible forecasts show total data-centre consumption more than doubling by 2030, with AI singled out as the key accelerator.

BY MICHAEL WANG, DC R&D DIRECTOR, AGINODE

AT RACK LEVEL, densities that were once confined to single digits are giving way to 10–30 kW as the new normal, while the most ambitious AI clusters push far beyond that and increasingly require liquid cooling as standard practice. These shifts are forcing utilities, operators, and planners to rethink grid connections, topologies, and specifications and physical placement of cabling and connectivity.

Balancing rising energy demands with efficiency gains

Over the past decade, the DC industry has made notable strides in energy efficiency. Average Power Usage Effectiveness (PUE) scores have significantly improved, declining from around 2.20 in 2010 to approximately 1.55 in 2022. Today, AI-focused data centres aim even lower, often targeting PUE values below 1.3. But despite these advancements, AI's rapid growth presents a paradox: while AI promises solutions for global sustainability challenges, the technology itself often significantly boosts power consumption.

AI workloads, particularly involving large-scale model training and inference, are exceptionally power-intensive, raising concerns about increasing energy consumption and carbon emissions. The International Energy Agency (IEA) projects that global electricity demand from data centres could double between 2022 and 2026, partly driven by AI. A University of Massachusetts Amherst study highlights the stark environmental cost of AI, noting that training a single large AI model could emit over 626,000 pounds (284,000kg) of CO_2. That's more than the lifetime emissions of five average cars.

AI is frequently positioned as a key to reducing global power consumption and promoting sustainability. However, realising these benefits currently demands increased energy and resource use. This tension embodies the 'rebound effect', where improved efficiency paradoxically leads to greater overall consumption. For instance, enhancing efficiency and lowering operational costs may inadvertently encourage more intensive AI usage, ultimately increasing total power consumption rather than reducing it.

The rapid growth of AI is reshaping R&D priorities and data centre design principles. The challenge facing the industry is clear: how do we harness AI's benefits without exacerbating environmental impacts? Power management and efficiency are becoming increasingly critical considerations. Addressing AI's inherent dichotomy – leveraging its transformative potential without escalating environmental consequences – is the paramount challenge facing data centre operations today.

Several strategies are already in development or deployment to tackle this challenge:

- **Renewable energy integration**: Transitioning data centres to renewable energy significantly reduces carbon footprints. Google's partnership with Kairos Power, for example, aims to deploy small modular nuclear reactors capable of delivering up to 500 megawatts of carbon-free electricity by 2030.
- **Optimised AI models:** Designing AI algorithms that effectively balance performance and energy efficiency can drastically cut down resource consumption.
- **Energy-Efficient hardware:** Investing in innovative, energy-efficient processors and smart integration systems reduces overall energy demands.
- **Advanced cooling techniques:** Employing cutting-edge cooling technologies lowers the energy needed to maintain ideal operational temperatures, enhancing overall efficiency.

Cabling and connectivity

With the recent surge in large-scale AI model launches, the need for massive-scale, high-performance infrastructure is becoming more urgent than ever. Network bandwidth and low latency have now become the primary bottlenecks in scaling AI data centers.

Large-scale AI models rely heavily on GPU interconnects, often using RDMA (Remote Direct Memory Access) to minimize CPU usage and boost access efficiency. These interconnects demand extreme bandwidth, and traditional Ethernet no longer suffices. While technologies like InfiniBand (a high-speed, low-latency network standard used in supercomputing) and RoCE (RDMA over Converged Ethernet, which enables fast data transfers over Ethernet) have improved overall networking, even 800G is becoming inadequate. The move to 1.6T is inevitable.

AI data center networks differ from traditional types in two key aspects:

- Demand for extremely high-speed networking.
- The need for ultra-low latency.

AI networks are leaping from 400G to 800G and on to 1.6T, which drives denser, standards-based structured cabling (ISO/IEC 11801-5, TIA-942) with MPO connectivity; higher per-channel speeds can cut cabling complexity and improve energy efficiency per Gbps.

Bandwidth improvement hinges on two key factors: the speed of each SerDes channel and the number of channels. (A SerDes channel is the end-to-end electrical path between two high-speed serializer/deserializer (SerDes) transceivers. In a data centre, that usually means the path from a switch/adapter ASIC's SerDes TX to a SerDes RX.) This largely dictates the max speed and reach you can run per lane (25G → 50G → 100G today, 200G emerging).

800G networks can be achieved via 16 x 50G channels, 8 x 100G channels of 4 x 200G channels. A 16-channel solution typically requires two MPO-16 (or MPO-24) connectors for full duplex communication. In contrast, a 4-channel solution needs only one MPO-8/12 connector, significantly simplifying infrastructure. The higher the individual channel speed, the lower the cabling cost and the better the energy efficiency per Gbps.

Latency is another major hurdle in AI performance. In large data centers, each additional switch layer increases latency. When scaling to 10,000+ GPUs, a three-tier switch architecture is often used, potentially adding up to five switch hops – each introducing latency that far exceeds that of the physical medium (fiber or copper) With growing scale and technical complexity, AI data center cabling must evolve faster than traditional data centers. A flexible, future-ready cabling infrastructure helps users adapt to rapid technological changes without repeatedly overhauling hidden infrastructure.

Standards such as ISO/IEC 11801-5 and TIA 942 emphasize how structured cabling can support ToR (Top-of-Rack), Spine-Leaf, and Mesh topologies—providing solid groundwork for dynamic and scalable architectures.

Connecting the AI-Ready Backbone: what to look for in Data Center Cabling?

Powering an AI-ready backbone starts with cabling that keeps pace with switch and optics roadmaps. In practice, that means native support for Base-8 and Base-16 parallel fiber so you can land today's 400G/800G SR links (e.g., 400GBASE-SR8 over MPO-16) and avoid stranded fibers on the way to higher speeds; most 800G pluggables ship with MPO-16 or dual-MPO-8/VSFF options, and the QSFP-DD/OSFP ecosystems formalize those interfaces.

Future-proof scalability also means choosing plant that's ready for the next Ethernet step: the IEEE P802.3dj project is defining 200/400/800 Gb/s and 1.6 Tb/s Ethernet, while industry forecasts show AI back-end ports moving rapidly from 800 G in 2025 toward 1.6 T by 2027 – so structured cabling, polarity methods, and connector choices should assume 100G- and 200G-per-lane optics.

Streamlined architecture and simplified installation are essential to reducing total cost of ownership (TCO). Pre-terminated, high-density systems have documented installation and risk reductions A cutting-edge solution could save up to 40% pathway space by consolidating dozens of patch cords into single assemblies with market leading small cable diameters.

High-density patch systems reduce cabling complexity and improve maintenance efficiency. Standards bodies are now formalizing very-small-form-factor (VSFF) connector families and kicking off multi-core fiber (MCF) work so selecting pathways and trays with headroom (bend radius, fill, labeling) helps you adopt these as they commercialize.

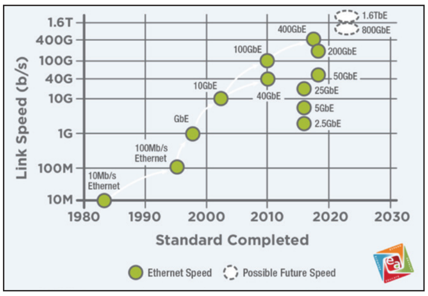

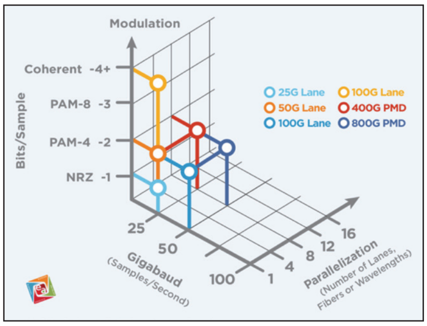

Channel count and 800G relationship (Source: Ethernet Alliance)

Across all sizes, from edge facility to hyperscaler, the recipe is consistent: design for density, engineer for migration, and instrument the physical layer so operations stay clean as speeds and port counts climb. Ensure you have the density and migration headroom (LC/MPO, HD/UHD), patching that supports an auditable workflow, and the option to right-size media by room and reach.

These foundations are tuned by site type. In enterprise/small private DCs, links are short and MACs (Move, Add, Change) frequent, so mixed copper/multimode plus UHD LC/MPO-LC modules keep front-of-rack neat while leaving a path to higher speeds. Colocation sites need fast turn-ups and clean tenant separation; HD/UHD frames, MPO backbones, strong labeling/shutters, and AIM-driven workflows reduce errors and speed audits.

Hyperscale/AI rows prioritize OS2 with Ultra low-loss MPO (Base-8/16), pre-terminated trunks, and micro-bundle cables to compress cabling volume at 100/400/800G – again with AIM to de-risk massive change rates. Edge/micro DCs and outdoor POPs are footprint-constrained and lightly staffed, so pre-terminated LC/MPO panels, remote visibility, and fit-for-purpose enclosures accelerate rollout and simplify upkeep.

Future faster with Elevate: driving data centres into the next generation

Elevate – A future-ready approach to data centre solutions

In today's data-driven economy, speed, resilience, and sustainability are the defining benchmarks for data centre infrastructure. Elevate – Future Faster, an Excel solution, was created with a single ambition: to help data centres and enterprise operators accelerate into the future with confidence.

Elevate brings together a comprehensive portfolio designed for ultra-dense, scalable deployments. At its core, Elevate provides:

- High-density fibre connectivity – precision-engineered for speed, reliability, and growth, including MPO solutions for high-fibre-count backbones and VSFF connector technology that delivers ultra-dense, space-saving performance in demanding white space environments
- Racks and containment – enabling airflow optimisation, modularity, and intelligent monitoring.
- Power distribution – iPDUs to modernise infrastructure sustainably.
- DCIM software – turning infrastructure data into actionable insights for operational efficiency and ESG reporting.
- Cooling solutions – from rear-door heat exchangers to hybrid liquid strategies for high-density loads.
- Fibre ducting and pathways – ensuring structured systems that simplify deployment while maintaining best-in-class performance.

Together, these solutions create a platform that allows operators to scale with confidence, balancing performance, sustainability, and cost-efficiency in environments where resilience and density are non-negotiable.

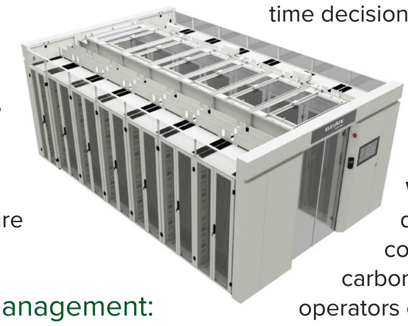

DCIM and power management: from insight to action

Data centres thrive on visibility. Without accurate, actionable information, operators struggle to keep pace with increasing density and sustainability requirements. That's why Elevate integrates intelligent power solutions with DCIM software to provide a single, holistic view of operational performance.

Elevate's partnership with Sunbird brings a DCIM platform capable of turning raw monitoring data into insights for optimisation and compliance. At the same time, nVent's high-density iPDUs deliver granular rack-level metrics on power and environmental conditions. Together, these tools create a digital twin of the data centre, enabling proactive capacity planning, risk management, and real-time decision-making.

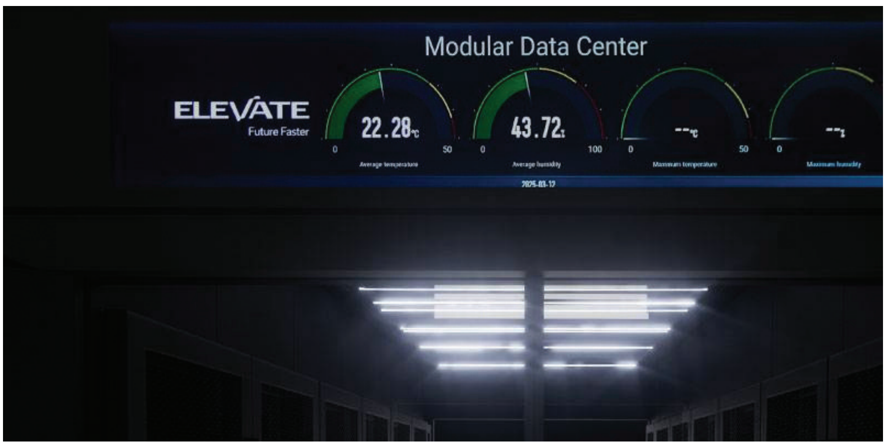

Crucially, this integration supports ESG and regulatory reporting. With transparent data on power consumption and carbon accountability, operators can meet compliance demands while also reducing waste and costs.

Racks and Containment: Building Blocks of Scalability

While visibility and control are critical, the physical backbone of any facility remains the rack and containment infrastructure. These systems dictate how effectively space, airflow, and cable management are optimised across the white space.

Elevate's Data Centre Rack (DCR) Series has been engineered specifically for ultra-high-density deployments in data centre and HPC environments. Supporting static loads of up to 2000kg, the racks combine strength with intelligent airflow management. Features include a vented front door and double-vented rear doors with up to 80% perforation, ensuring effective cooling even under heavy load. Fully adjustable 19" profiles are fitted with airflow management baffles that move with the rails, directing cold air precisely through active equipment.

The DCR also addresses operational efficiency with overlapping brush panels, blanking options, integrated cable management, and high-volume cutouts for cable entry. Quick-mount PDU trays and extended roof configurations further support deployment flexibility, while 64A Commando plug clearance makes power distribution easier to integrate.

When paired with Elevate's hot and cold aisle containment systems, the DCR Series enhances thermal performance across both white and grey space.

By reducing bypass airflow and improving cooling efficiency, the solution helps operators maximise density while lowering energy costs.

Partnerships Driving Innovation

Elevate's strength lies not just in its solutions, but in its ecosystem of partners, carefully selected to complement and enhance its portfolio. This collaborative approach ensures customers gain access to a joined-up platform of innovation across the data centre lifecycle:

- **Sunbird –** transforming DCIM from monitoring to actionable intelligence.
- **Senko –** providing precision fibre connectivity for next-generation networks.
- **nVent –** delivering cooling and intelligent power distribution.
- **Schleifenbauer –** delivering intelligent power distribution.**[RB1]**
- **Axis, Avigilon and Suprema –** enabling layered security from perimeter to rack.

These partnerships mean customers don't just deploy hardware; they benefit from a future-ready ecosystem that adapts as demands evolve.

Events and Exhibitions: Connecting with the Industry

Elevate's role in the market extends beyond technology - it is actively engaging with customers and partners across Europe and beyond.

Alongside this flagship event, Elevate will also be exhibiting at key European shows throughout the rest of the year:

- **DCW Madrid –** 16 – 17 October
- **DCW Paris –** 5 – 6 November

These events provide opportunities for customers to connect directly with Elevate's experts, see solutions in action, and understand how the brand is helping operators design for tomorrow's challenges.

Elevate's RePower Trade-In for Tomorrow* initiative strengthens this proposition further. Customers can receive up to £35 per legacy PDU traded in, which can be offset against the purchase of new, intelligent, energy-efficient models. This provides a clear financial incentive to modernise ageing infrastructure, making upgrades more cost-effective and accessible.

But the value of RePower extends beyond cost savings. Many legacy PDUs are large, inefficient, and carbon-heavy in operation. Through RePower, these units are removed, recycled, and disposed of responsibly, preventing unnecessary e-waste and ensuring materials are processed in line with environmental best practice. By replacing outdated hardware with intelligent PDUs, customers not only reduce their operational carbon footprint but also benefit from enhanced monitoring data that feeds directly into DCIM dashboards.

In this way, RePower is both a sustainability initiative and a management enabler, reducing environmental impact while unlocking smarter visibility, richer insights, and better reporting across the white and grey space.

* Terms and conditions apply. Contact the Elevate team at elevate@excel-networking.com or call 0121 326 2471 for further information.

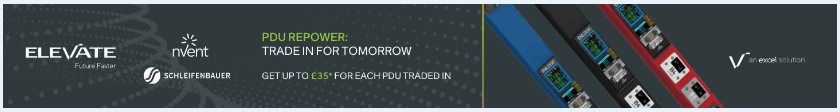

Conclusion: Elevating the Future of Data Centres

As the data centre industry accelerates towards higher density, greater accountability, and sustainable growth, Elevate provides the platform for performance that operators need.

By combining intelligent power and DCIM software, scalable racks and containment, and a partner ecosystem delivering cooling, connectivity, and security solutions, Elevate is enabling facilities to move future faster.

With initiatives like RePower, a packed calendar of industry exhibitions, and its upcoming Pole Position event, Elevate is not only delivering innovation but also shaping the conversation around the future of passive infrastructure.
Discover Elevate - Future Faster.

"Elevate is about delivering speed, precision, and innovation - but also trust. Backed by the combined strength of Sonepar, Mayflex and Excel Networking, customers know that they're investing in solutions that are future-ready and fully supported."
Andrew Percival, Managing Director

"DCIM is no longer a 'nice to have' — it's mission critical. By linking intelligent power hardware with analytics and automation, Elevate helps customers manage density, reduce waste, and meet regulatory requirements head-on."
Simon Jacobs, Product Manager

This October, Elevate will host its own exclusive end-user event:

- **Pole Position:** Data Centres in the Fast Lane
- **Date: Thursday** 9 October
- **Venue:** F1 Arcade, St Paul's, London
- **Time:** 12.00pm – 6.00pm

The day will feature presentations from Elevate and its partners including nVent, Sunbird, Senko, Axis, and RED Engineering, covering the future of cooling, DCIM, fibre connectivity, and DC design. Attendees will then shift gears for an afternoon of F1 simulator racing, networking, and prizes - including an Elevate-branded iPad Air, Apple HomePod, and AirPods.
Secure your spot on the starting grid - because the future doesn't wait.
Register to attend

Shaping the future of coolant technology

In the fast-changing world of automotive engineering, data centre infrastructure, and industrial processes, the importance of thermal management is often underestimated. Yet, behind every vehicle engine, every electric vehicle battery, every AI-driven data centre, and countless industrial applications, a critical enabler technology needs to be carefully considered: How to manage the excess of heat and cool the installations for optimal performance. One company that has quietly built a global reputation for excellence in this field of expertise is Arteco.

ARTECO was founded in 1998 as a joint venture between Texaco (now part of Chevron) and Elf (now part of TotalEnergies). Both parent companies brought decades of know-how in coolant technology as far back as 1948, with Texaco having pioneered breakthrough coolant OAT technology already in 1985.

In 1992, Texaco introduced in the market his groundbreaking OAT (Organic Additive Technology) coolant, a major advancement in engine coolant corrosion protection. Today, it is still the backbone for the modern quality coolant technologies. This pioneering spirit would become part of Arteco's DNA.

In December 2016, Arteco transitioned into a standalone company while maintaining a 50/50 joint venture structure between Chevron and TotalEnergies. Headquartered in Ghent, Belgium, with a production facility near Antwerp, Arteco has grown into a truly global player.

A recent milestone for Arteco was the opening of their new production plant in Nantong, China, reflecting the company's ambition to serve the fast-growing Asian market and allows them to enforce their capabilities in Asia. Arteco is positioning itself as a global leader in advanced coolant solutions. Technology at the core.

If there's one defining characteristic of Arteco, it's the emphasis on technology and innovation. Nearly a third of the company's workforce is employed in technology, underscoring the strategic importance of technical leadership towards their customers.

At its R&D centre in Ghent, Arteco develops advanced coolants and heat transfer fluids tailored to diverse applications such as automotive and New Energy Vehicles, Data Centres and Heat Transfer fluids for industrial applications.

A new laboratory in China will further expand this capability, reflecting the company's commitment to being close to its customers in key markets.

Reinventing data centre cooling

While Arteco has deep roots in the automotive industry, one of its most exciting frontiers is data centre cooling. The rise of artificial intelligence and machine learning has turbocharged thermal management demand for high-performance computing, pushing traditional air-cooling systems to their limits. As processors consume more power and corresponding generate more heat, liquid cooling is emerging as THE viable solution to go to.

Arteco is seizing this opportunity by leveraging decades of coolant expertise to design products for direct-to-chip applications.

In 2023, it launched the ZITREC® EC range, specifically tailored for data centre environments.

This portfolio includes:

- **ZITREC® EC 10:** A water-based, biostatic coolant that resists biofouling without the use of biocides.
- **ZITREC® EC 20:** A baseline propylene glycol solution offering proven reliability.
- **ZITREC® EC 30:** An ethylene glycol-based coolant with a flux compensation package, designed to protect complex aluminium radiators from corrosion.
- **ZITREC® EC 40:** A low-conductivity formulation.

What sets Arteco apart is not just the products themselves, but the support ecosystem surrounding them. Many data centre operators and server manufacturers are familiar with IT infrastructure but not with coolants.

Arteco bridges this gap, offering expertise in compatibility testing, regulatory compliance, waste management, and product stewardship. By doing so, the company positions itself not merely as a supplier but as a strategic partner in the data centre ecosystem.

Sustainability – Accelerating Change Together

Perhaps the most defining feature of Arteco's current strategy is its commitment to sustainability.
Earlier this year, the company refreshed its sustainability strategy Arteco ACT, which stands for: "Accelerate Change Together."

The strategy is built around 3 key focus areas:

Climate Action

- Net-zero emissions for Arteco's own operations by 2030.
- Net-zero across its entire supply chain by 2050.

Smart Use of Resources

- Reducing reliance on fossil-based materials.
- Moving towards a circular economy.

People and Community Care

- Fostering a diverse, inclusive, and safe workplace.
- Supporting local communities through volunteer initiatives and social engagement

The ECO product range

A tangible outcome of Arteco ACT is the forthcoming ECO product range, which will be rolled out across all segments, from automotive to data centres and industrial fluids. The process was rigorous.

Arteco conducted life cycle assessments (LCAs) of its major products, not only examining carbon footprint but 16 different Product Carbon Footprint environmental impact categories.

The analysis revealed that the two biggest contributors to environmental impact were raw material sourcing, particularly fossil-based base fluids (PG and EG) and the end-of-life phase of coolants.

Arteco's ECO coolants incorporate base fluids originating from renewable feedstocks or recycled sources using a mass balance approach to allow for traceability and certification. These base fluids, Monoethylene Glycol (MEG) or Monopropylene Glycol (MPG) are traditionally virgin grade material derived from fossil resources.

Our ECO products help decrease reliance on fossil resources and have a significantly lower Product Carbon Footprint (PCF) compared to the virgin fossil product.

To confirm the traceability and reliability of this process, Arteco has obtained the International Sustainability and Carbon Certification (ISCC) PLUS for its mass balance approach towards more sustainable resources.

Crucially, customers were involved throughout the process. They demanded ECO products that would not compromise application performance.

This is in line with Arteco's vision on quality.

The ECO range also reflects a pragmatic approach to adoption as it is a drop-in solution, ready to use immediately.

Beyond the mass balance approach, Arteco envisions a circular future for coolant use.

To support customers, Arteco provides full data transparency, offering product carbon footprints and regulatory support through its product stewardship team. When it comes to sustainability, robust data and transparency will be key as to ensure trust. The data centre industry in particular faces a paradox: exponential growth in energy consumption driven by AI, alongside ambitious net-zero commitments.

Coolants may not be the largest contributor to energy use, but as Arteco notes, they represent 'low-hanging fruit'. Switching to ECO-based coolants can offer an immediate, one-to-one reduction in product carbon footprint, while maintaining performance and quality without disrupting operations. For customers under pressure to demonstrate tangible progress, this is a simple yet effective lever.

Of course, challenges remain. ECO products are currently more expensive than traditional alternatives, and volumes of sustainable glycols are still limited.

Political shifts can also temporarily slow momentum on sustainability. But Arteco remains optimistic – while Arteco may only be a droplet on the plate, a lot of droplets can really cool the plate. This philosophy captures the company's belief that collective action, however small each contribution may seem, is essential in addressing climate change.

Conclusion

To support adoption of its technology innovations, Arteco is investing in customer education. On November 27th, the company will host its ECO Webinar, an in-depth session exploring the ECO product range, methodologies, and implementation strategies.

The event will feature technical deep dives from Arteco experts, followed by Q&A. Customers will be able to register via the company's website and LinkedIn page.

With a relentless focus on technology, a clear sustainability roadmap, and an openness to collaboration, Arteco is carving out a leadership role in both established and emerging markets.

From cooling traditional engines to enabling the next generation of data centres, Arteco is proving that innovation and responsibility can go hand in hand. Its journey offers valuable lessons for other companies – that agility, technical excellence, and a values-driven approach can create impact far beyond the company's size.

As the world grapples with climate change and the demands of digital transformation, companies like Arteco can do their part in ensuring that progress is not only possible, but sustainable.


R&M
BLADESHELTER
R&M

Light in the data centre: Efficiency meets future-proofing

Why intelligent lighting solutions matter today - and set tomorrow's standards. Evolving demands - lighting has long been part of the solution

BY TRILUX

THE ENERGY REQUIREMENTS of digital infrastructure are constantly increasing. Data centres are among the most energy-intensive consumers in modern societies. While server technology, cooling and building systems have made significant progress in recent years, the focus is now shifting to lighting. It directly impacts not only energy use but also thermal load – and, in turn, the overall energy efficiency of the entire system.

"A key factor is replacing conventional fluorescent tubes with modern LED technology, which is up to 85% more efficient, generates less heat, and lasts significantly longer- a clear contribution to energy savings and reduced maintenance costs," explains Claudia Lüdenbach, Head of Application Management Industry Marketing at TRILUX, speaking at the Data Centre World Frankfurt trade fair, adding: "But efficiency alone is not enough. Smart control, sustainable materials and design, and a good working environment for staff are equally important."

Light in the server room: precise, efficient and reliable

Good lighting is crucial for accurate work in the aisles between the server racks - whether for maintaining components or reading displays. High vertical illuminance, glare-free light and excellent colour rendering (CRI > 90) are essential. At the same time, the lighting system must produce minimal heat to reduce strain on the air conditioning system.

As a pioneering technology and service provider in the European lighting market, TRILUX offers LED luminaires with efficient thermal management and high dust resistance to meet these requirements. Modular continuous line systems like E-Line Pro optimally light server aisles and surrounding areas.

LiveLink lighting management - At a glance

- Helps to save up to 85 % energy through demand-driven control
- Seamless integration into building management systems
- Can be combined with sensors, scheduling and zone control
- Time-saving commissioning
- Intelligent functions

Smart Lighting Solutions for Every Area

TRILUX ensures excellent lighting not just in server rooms, but throughout the entire data centre – including offices, technical areas, outdoor spaces, and circulation zones. Each has unique lighting needs, and TRILUX provides tailored solutions that combine comfort, safety, and energy efficiency - all from a single source. Intelligent control systems allow the lighting to be regulated as required - for optimum conditions and low operating costs throughout the building complex.

Office Areas

Create optimal work environments with smart lighting controls that adapt to user needs while reducing energy use.

Plant Rooms

The building technology in data centres is designed for maximum reliability. TRILUX supplies robust, higher protection rated lighting solutions to meet the demands.

Circulation Areas

TRILUX lighting solutions help to create easy orientation in corridors and staircases and a feel-good atmosphere in break rooms.

The Building Surround

Lighting for entrances, façades, walkways, and car parks enhances safety and integrates the facility into its urban surroundings.

An integrated lighting management system such as LiveLink can network all these areas - indoor and outdoor -controlling them centrally and improving overall building efficiency and quality.

Smart positioning and light control create a uniform, eye-friendly light distribution that increases safety and minimises the risk of errors during maintenance work.

Focus on sustainability - a cycle not a short life cycle

Sustainability in data centres means much more than just reducing CO_2. It is also about the service life, reparability and recyclability of technical components. That's why TRILUX develop durable, long-life lighting systems that can be adapted to future requirements - for example, with retrofittable modules or intelligent control systems.

Reusable housings, modular construction and service-friendly designs already make many modern lighting solutions fit for the circular economy. They also fulfil the requirements for sustainability certifications such as LEED or BREEAM - increasingly important for operators. The economic benefits are also clear: low-maintenance products reduce operational costs and minimise downtime, improving efficiency and offering greater planning security.

Continuous operation requires absolute reliability

Data centres operate 24/7 - and so must their lighting. To minimise maintenance costs and failure risks, operators rely on durable luminaires with an integrated emergency lighting function and high electromagnetic compatibility. All systems must work reliably even under extreme conditions - this applies to lighting as well as control and integration into the building automation system.

“Modular architectures and robust materials allow systems to be flexibly adapted and expanded as required. This not only creates technological security, but also investment protection. A lighting solution that grows with you offers long-term added value,” summarises Claudia Lüdenbach, Head of Application Management Industry Marketing at TRILUX.

Expertise, planning skills and international experience are particularly crucial in data centre construction. Lighting technology and service providers such as TRILUX, who have set standards in large-scale industrial projects, contribute their expertise to lighting design for data centres - at every stage: from needs analysis to lighting design and implementation.

Thanks to its global presence and close networking with planners and operators, TRILUX develops customised concepts that are economical, scalable and future-proof. Efficiency, safety, and service always take centre stage, which is complemented by the flexibility to react quickly to new requirements. Visitors to the “Data Centre World Frankfurt” trade fair will be able to discuss lighting systems in data centres in more detail with the lighting experts at the TRILUX stand (H063 | Hall 8) on 4-5 June 2025. www.trilux.com/data-center

Lighting solutions with vision

Future-proof lighting for data centres combines efficiency with intelligent control and sustainability with adaptability - delivering reliability in complex environments. Whether it's LED technology, modular continuous line lighting systems or networked lighting management, planning for today means preparing for tomorrow.

Simplify Your Light - also in the data centre.

Unlocking bandwidth, connectivity, and power for AI data centres – with next-level efficiency

The AI era – with its GPU/CPU-intensive workloads—is exposing the limits of traditional data centre design. Training clusters now connect tens of thousands of accelerators and drive unprecedented east-west traffic.

BY CARSTEN LUDWIG, MARKET MANAGER DATA CENTRE, R&M

INTER-SITE REPLICATION is running at hundreds of gigabits per wavelength—and rising. To keep pace, data centres must evolve well beyond legacy assumptions in processing, power, cooling, and connectivity. The good news: with the right investment strategy, physical infrastructure, and network design, operators can turn these pressures into durable advantages.
Investment efficiency: flexibility over one-off bets

Bandwidth, port density, and power technologies must scale quickly, but not at the cost of over-specifying or overspending. Building credible business cases means:

- Meeting current needs and foreseeable growth with acceptable time-to-ROI,
- Avoiding vendor lock-in through standards-based designs, and
- Supporting multiple platforms with modular, reconfigurable building blocks.

Rapid densification (100 kW racks, megawatt rows, 576-GPU deployments) and the fact that no single design is the 'winner' were highlighted at the recent Data Centre EXPO at the RAI Amsterdam. The implication: prioritise portfolio thinking and design flexibility to manage capex/opex and risk. Architects should optimise from the outset and keep room to adjust footprint and topology as requirements change.

Power and cooling efficiency: design for "performance per watt"

During the Data Centre EXPO, the industry was urged to curb over-provisioning and make performance per watt the north star, rather than raw output. Several recommendations materially reshape lifetime TCO and should be designed-in from day one:

- **Beyond "net zero" to Absolute Zero**: real-time energy transparency, circularity, measurable outcomes—and a hard stance on over-provisioning.
- **"Utility Net-Zero" architectures:** on-site generation, hydrogen fuel cells for long-duration backup, microgrids, and bidirectional grid services to improve electrical efficiency, resilience, and emissions intensity while progressing toward 24/7 clean operations.
- **Water-smart, heat-savvy cooling:** liquid-immersion cooling (eliminating process water), rigorous heat reuse into district networks, and CO_2-based distribution loops that turn waste heat into a resource and shrink cooling energy.

The Ethernet Alliance expects AI to push data-centre electricity use sharply higher this decade, making energy per bit a first-order network parameter. Every watt saved in optics and cabling is either returned to workload or never drawn from a constrained grid. One practical lever: IP-over-DWDM. Case studies show that moving to router-based coherent optics with open line systems can deliver major energy savings (reported up to 97% in certain provider networks) alongside capex/opex reductions from fewer boxes and less space. With an instrumented physical layer, AI-driven operations can then optimise subsystems such as cooling in real time – often cutting energy consumption by 20-40%.

Conventional air cooling alone is no longer sufficient for high-density GPU configurations. In hotter, water-constrained regions (e.g., parts of Spain and Portugal), cooling complexity and cost rise further; a typical water-cooled facility can consume 1-2 litres of water per kWh depending on design. Permitting and environmental impact processes can be slow, sometimes face local opposition, and raise development risk. All of this reinforces the value of efficient, low-water cooling and careful siting.

Connectivity efficiency: build for density, observability, and smooth upgrades

To handle exploding data volumes and tight latency budgets, operators must extract the maximum from their fibre plant – both in white space and across the metro. That calls for:

- Higher port densities and migration to very small form factor (VSFF) connectors.
- Reduced-diameter ribbon fibres and high-density platforms with integrated asset/environmental monitoring.
- Smart migration paths so today's choices don't block tomorrow's 800 G and 1.6 T rollouts.
- Integrated power and liquid/hybrid cooling for 30 kW+ racks and beyond.
- Robust racks/ODFs maintaining bend radius compliance and cable management as density climbs.

If the passive layer is well specified and perfectly integrated, you can then layer in AI/DCIM to run the plant more efficiently. Dense WDM (DWDM) is central to this, multiplying capacity on existing fibre by stacking many wavelengths. Coherent optics dramatically increase per-wavelength throughput and reach with greater impairment tolerance, while DWDM readily carries both coherent and legacy on-off keying signals side by side.

Coherent pluggables–400ZR now, 800ZR/ZR+ emerging—plug directly into routers and switches, enabling IP-over-DWDM. This collapses footprint and power, lowers cost per bit, accelerates turn-up, and unlocks more capacity per fibre pair while providing a clean path from 400 G to 800 G per wavelength. QSFP-DD/OSFP modules over single-mode pairs deliver 400/800 Gb/s per lambda, with single-mode increasingly popular for DCI and long fabric runs. Hyperscalers are scaling shipments of 400/800 G coherent optics and adopting 800ZR+ to stitch together distributed AI campuses where land and power constraints force multi-site architectures.

A pragmatic migration path

- **Design for 800 G now, 1.6 next.** Choose fibre/cassette systems and ODFs that support higher-lane optics and coherent pluggables without re-terminations.
- **Adopt IP-over-DWDM for metro DCI.** Standardise on 400ZR/ZR+ with a roadmap to 800ZR+, leveraging open line systems to scale wavelengths without vendor silos.
- **Instrument everything.** Specify fibre infrastructure with asset tracking and environmental sensors so AI/DCIM systems can "see" the plant and optimise it continuously.

Why repeatability matters

Approaches that can be repeated and scaled are essential for AI growth. Prefabricated connectivity (modules, cassettes, cabinets) reduces deployment time, cost, and complexity – enabling predictable rollouts. Treat repeatability as a best practice for managing physical connectivity changes (moves/adds/changes) through DCIM/AIM. Fibre platforms should be simple to change, upgrade, or scale in just a few steps. A zoned network architecture (End-of-Row / Middle-of-Row) is a good example: a repeatable building block you can deploy across rows to reduce devices and improve port utilisation.

An integrated approach: quality products plus holistic design

High-quality components are essential – but quality alone isn't enough. Using parts that meet stringent standards and are tested for performance and availability keeps networks future-proof and investments secure. Equally important is structured, forward-looking design and documentation. In short: pair top-tier connectivity hardware with holistic planning; otherwise, excellent parts can still produce faults or downtime.

Cabling quality is critical to uptime; a large share of IT failures originate in the cabling system. Measurement and testing are key – from planning through formal acceptance. Manufacturers should offer comprehensive QC and field-testing support, and products should guarantee permanently stable transmission quality.

A complete solution also considers access devices such as carefully selected SFPs and coherent modules to secure high-quality signal transfer. As speeds rise (400/800 G and beyond), transceiver formats (QSFP-DD/OSFP, MPO-16, etc.) drive cabling and connectivity choices. Optics and physical layer must therefore be designed together from the outset, not treated as an afterthought.

Signal quality at the optics interface is pivotal: low insertion loss and high return loss keep transceiver links clean and stable. Every dB saved on the passive side increases transceiver headroom and reduces flaps/BER. That makes link-loss budgets a foundational design consideration: how much loss can you allocate to fibre versus connections? The push to deploy AI closer to users (edge sites) adds environmental challenges – vibration, dust, EMI, limited space – which require rugged enclosures, efficient small-footprint modules, and remote monitoring.

Redundancy, utilisation, and loss-budget discipline

- **Engineer redundancy to the target Tier.** For higher tiers, implement dual A/B paths with N+1 (or more) at the layers that matter – power, cooling, and connectivity. That includes dual cross-connects, diverse routes, separate panels/ODFs, and diverse meet-me rooms. Tier III requires redundant capacity; Tier IV is fully fault-tolerant. Choose the model explicitly and build cabling to match.
- **Document and monitor.** Manage redundancy through DCIM/asset tracking so every path is visible, documented, and alarmed.
- **Respect attenuation budgets.** Every connector and splice "spends" loss budget. Minimise mated pairs and hops in each optical path. Where patching is unavoidable, use low-IL grade cassettes/cords and keep polarity consistent to avoid extra adapters.
- **Optimise utilisation.** Architect DCIM and conversion methods for maximum trunk-fibre utilisation (no dark strands stranded), and use optimised breakout/aggregation to avoid wasting fibres or lanes.

Designing, building, and scaling for AI efficiency

AI workloads, rising power densities, and sustainability pressures are reshaping the data-centre landscape. Successful operators now design the full stack – compute, cabling, power, cooling, room layout – as a single, scalable system. No component is planned in isolation.

AI-scale performance depends on a meticulously engineered optical and power stack: dense, clean, well-managed fibre; DWDM with coherent pluggables to extract multi-terabit capacity from existing pairs; and an instrumented layer that lets AI tune cooling, power, and traffic hour by hour. Converging routing and optics with IP-over-DWDM trims space, watts, and cost per bit, while real-time telemetry enables proactive operations and measurable efficiency gains. Paired with advanced liquid/hybrid cooling and a performance-per-watt mindset, this approach turns bandwidth, connectivity, and power from constraints into competitive advantages – resilient across geographies, aligned with tightening ESG expectations, and ready for the next wave of AI growth.

In the words of Susanna Kass (Digital Gravity Infrastructure Partners): "There is no AI without data centres. 'AI factories' are a special breed of data centres that focus on efficiency and performance, and energy assets integrated with AI technologies. They are critical infrastructure for Digital Sovereignty to increase economics efficiency, humanity safety and innovation."