



DATA CENTRE SOLUTIONS

ISSUE | 2026

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

[DATACENTRE.SOLUTIONS.COM](https://datacentre.solutions.com)



RETHINKING AI INFRASTRUCTURE

the lifeblood of every modern data centre



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

scan to explore



A new year, but same old problems?

➤ SALUTE'S new State of the Industry report reveals that the AI boom is pushing data centre infrastructure to its limits, with 83% of senior leaders warning that AI workloads will exceed their current capacity within two years. Time to panic?

I remember being at an industry event 18 months or so ago when an analyst showed a slide with anticipated data centre capacity disappearing off the top of the graph, while the power availability line 'bumbled' along the bottom of the chart, with only a slow, steady increase. The gap between the two was huge and getting bigger...thanks to AI.

And a more recent conversation with a UK power provider revealed that, if all the data centres in their region come to fruition, the power demand will more than double by 2030 or so.

Add in the slightly less high-profile problem of recruiting enough skilled personnel to the data centre sector to meet this huge growth demand and one wonders why no one has hit the panic button just yet.

Then again, the money markets (and who would ever doubt them?!) seem to spend half their time convinced that we are witnessing an AI bubble, and that there will be a huge adjustment at some stage, and, therefore, there will be far fewer new data centres required than the current 'hype' suggests.

And I have had conversations with other energy suppliers who take a similar view – once the many speculative projects and competing ones are whittled down as the

number of (AI) customers is far less than anticipated, then there's every chance that power suppliers are well positioned to meet a far less spectacular data centre growth over the next few years.

And then there's AI at the edge. No need for the huge, LLM data centres, but many, many local ones to deliver the required AI application performance close to the users. Fascinating times for sure!

While I won't claim that DCS has all the answers, we are starting a series of regional Roadshows, with networking and collaboration very much the focus, designed to understand what's going on in the data centre world in some of the emerging markets. We start in Newcastle later this month. We are visiting London in March – with the focus very much on how the data centre industry is 'spreading its wings' further afield from Slough and Docklands, before moving on to Wales and Scotland. You can find more details on the DCS website.

Finally, I'm delighted to recommend to you our expanded DCA content in this, and subsequent, issues of the magazine. At the back of the magazine you will find the customary content - the welcome from Steve Hone and some technology articles. Earlier on in the magazine, you'll find some Industry Insights – focusing on topical issues. Finally, finally, we'll be at DCW in early March, where Mark and I will be out and about on the show floor and also be based on the DCA stand, conducting video interviews.

We look forward to seeing you there!

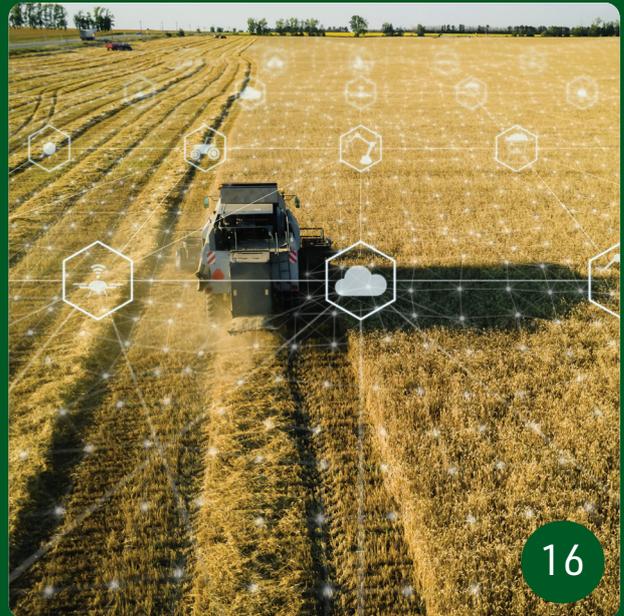


Contents

Cover Story

Rethinking AI Infrastructure

AI and data centres are not neutral tools - they are embedded in systems that affect water, energy, climate and equity. We cannot afford to leave those systems unexamined or unaccountable.



16

18 Can DCIM's AI-driven future increase data centre sustainability?

The rapid expansion of AI, big data, and IoT has intensified the complexity of data center capacity and infrastructure management.

22 Partner data centres with biodiversity for a sustainable AI era

Do we want to be remembered as an industry that opposed our natural world and caused extinctions, or one that set an example of industry working alongside nature?

30 Redefining the journey: tech-agnostic microgrids for data centres

At AVK, we believe resilience comes from choice. No single technology can power tomorrow's data alone.

35 Future-proofing data centres through on-site power and distributed carbon capture

There is finally a pathway for data centres to address both power needs and decarbonization, while maintaining full control over power supply.

36 Beyond cooling: how data centres can become heat hubs for Net Zero infrastructure

Every year, data centres in London release enough waste heat to warm up to half a million homes.

38 Unhackable by design: securing AI data centres at the physical layer

AI data centres are fast becoming the backbone of the digital economy.

40 The firewall: mitigating data centre fires with stone wool sandwich panels

In a time when digital continuity directly impacts public safety, economic stability, and social interaction.

44 The five Ss: a guide to selecting effective data centre fire suppression

The modern data centre may not always be a single, large space and may instead be a series of data halls.

48 How to get your physical data centre security right from the start

Data centres are, rightly, renowned for their physical as well as digital security.

56 The greening of the Internet Exchange Point

Following the recent COP30, the conversation around digital sustainability has never been more urgent.

58 Green data centres: how to turn sustainability into strategic value

Over the last decade, we've seen data centre demand grow rapidly to support the rise of Artificial Intelligence (AI), cloud and hyperscaler workloads.

60 The lifeblood of every modern data centre

The complex data centre supplier landscape, is experiencing unprecedented growth with new suppliers entering the DC market every day, what are the real differentiators?

NEWS

06 Data centre outlook 2026: growth and trends

07 New report compares big tech's approach to nature in data centre plans

08 Vertiv report highlights emerging AI data centre trends

09 EMEA data centre investment on the rise

10 Data centre programme targets AI-focused startups in Hong Kong

11 Strategic joint venture develops AI infrastructure in Saudi Arabia

12 New sustainable data centres under construction in Spain

13 Accenture strengthens AI data centre services with DLB acquisition

14 Data centre outsourcing market size to cross \$243.3 billion by 2034

15 Innovative solar project marks new era for UK data centres



DCS DATA CENTRE SOLUTIONS

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

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

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

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

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

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

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

Directors
Scott Adams: CTO
Sukhi Bhadal: CEO

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

Angel BUSINESS COMMUNICATIONS

Datacentre Solutions is published 10 times a year on a controlled circulation basis in Europe, Middle East and Africa only. Subscription rates on request. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd. © Copyright 2026. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. ISSN 2756-1143 (Online)

Data centre outlook 2026: growth and trends

Data centres are set for expansion, driven by AI and sustainability transformations, with global investments enhancing infrastructure strategies and energy solutions.

THE LANDSCAPE OF global data centres is gearing up for an expansion, led by advancements in artificial intelligence (AI) and energy strategies. According to JLL's latest Global Data Center Outlook, capacity is anticipated to nearly double from 103 GW to 200 GW by 2030. With AI expected to account for half of this capacity, the sector is poised for changes.

The projected growth seeks investment, estimated at \$3 trillion over the next half-decade. This figure includes \$1.2 trillion aimed at real estate asset creation and around \$870 billion in new debt financing. Market authorities are observing a shift akin to the original cloud migration of the past.

AI's Pervasive Influence

AI's integration is altering the nature of data centres. By 2030, AI workloads could represent 50% of data centre capacity, from 25% in 2025. JLL predicts an important moment in 2027 when AI inference workloads surpass

training requirements. AI training facilities now require higher power densities, outpacing traditional data centres in lease rate premiums.

The development of speciality AI chips is predicted to strengthen their revenue share from 20% to 50% in the semiconductor market by decade's end. Emerging technologies, such as neuromorphic computing and ultra-efficient task processing, potentially curbing infrastructural demand while promoting energy efficiency.

Growth Across Regions

In regional dynamics, the Americas will maintain dominance, holding 50% global capacity with the fastest growth rate. Asia-Pacific is forecasted to grow from 32 GW to 57 GW, with Europe, the Middle East, and Africa (EMEA) set to expand by 13 GW, driven largely by demand in established European and emerging Middle Eastern markets.

While APAC's growth is led by colocation, there's a decline in on-premise capacities. For EMEA, demands from hyperscalers in key hubs, alongside sustainability initiatives, are guiding developments.

Energy and Sustainability Efforts

Efficiency and sustainability challenges remain significant, with grid connection lead times challenging operators. Many are investing in self-generated energy to offset utility delays, echoing initiatives seen in places like Dublin and Texas. Innovations in battery energy storage systems (BESS) are gaining traction, aiming to manage short-term outages efficiently and hasten interconnection timelines.

The push towards renewables continues, with initiatives combining renewable power sources to cut costs and environmental impact. Hyperscalers already counterbalance their U.S. portfolios with renewable energy, while EMEA projects aim to integrate private transmission systems to enhance power affordability for tenants.

As expectations around renewable energy intensify, nuclear power is emerging as a potential reliable option, though significant capacity is likely several years away.

Capital Market Progressions

The capital landscape is maturing concurrently, with fundraising strategies now prioritising core investments. A shift towards recapitalisations and joint ventures is evident, particularly with the rise of AI and neocloud technologies.

Fresh financial strategies are evolving to support such modern ventures, ensuring frameworks accommodate the capital inflows while safeguarding investments and nurturing potential growth.



New report compares big tech's approach to nature in data centre plans

Calls for data centres to, at a minimum, restore enough nature to compensate for operational emissions.

A NEW REPORT, thought to be the first scorecard to compare nature strategy across Europe's data centres, reveals striking differences in how the world's largest technology companies integrate nature into their data centres.

The Europe's data centres: The nature report card, from leading European carbon-removal developer Arbonics, reviews Google, Microsoft, Meta and Apple's European data centres. It compares and ranks their approaches to land use, site design and ecosystem restoration, and evaluates how much nature restoration would be needed to balance each facility's annual emissions.

As demand for AI accelerates, the footprint of Europe's data centre sector faces growing scrutiny. Electricity demand is projected to rise from 96 TWh in 2024 to 168 TWh by 2030 and 236 TWh by 2035, a 150% increase in just over a decade. The International Energy Agency expects Europe to remain one of the world's largest regions for data centre electricity use.

The companies driving the AI race are also, in many cases, at the forefront of integrating nature into their operations. Microsoft emerges as the report's overall leader, with 6,414 hectares of land permanently protected (more than its estimated global data centre footprint) and over 77,000 trees planted through community projects. Its initiatives include native planting at Middenmeer in the Netherlands, large-scale tree planting in West Dublin, and riverbank restoration in Spain.

However, the report warns that despite encouraging examples, nature-first design is far from standard practice. The development of large hyperscale facilities continues to involve land conversion, material-intensive



construction, and changes to local ecosystems. Arbonics argues that because of this, operators should go further and restore enough nature to compensate for their annual operational emissions.

Forests offer a clear pathway to repairing the pressures linked to data centre development. Trees sequester carbon, rebuild soils, regulate water and support biodiversity. One year of operations at Meta's Luleå data centre in Sweden would require restoring 3,350 hectares of forest, equivalent to planting 8.4 million trees, according to the report. Google's Hamina site in Finland would require around 19.4 million trees — enough to cover the city of Paris.

"Data centre operators can help re-establish the ecosystem processes that support their infrastructure, creating long-term ecological value alongside their climate commitments" said Lisett Luik, co-founder of Arbonics.

The report acknowledges that past industrial growth cost Europe much of its nature; forests once covered roughly 80% of the continent, falling to less than half by the end of the seventeenth century. The next wave of growth, driven by AI, can take a different path. To help operators act on this, it outlines

four priorities:

- Increase land restoration
- Prioritise brownfield over greenfield development
- Report biodiversity at site level
- Integrate nature-led design features, such as green roofs

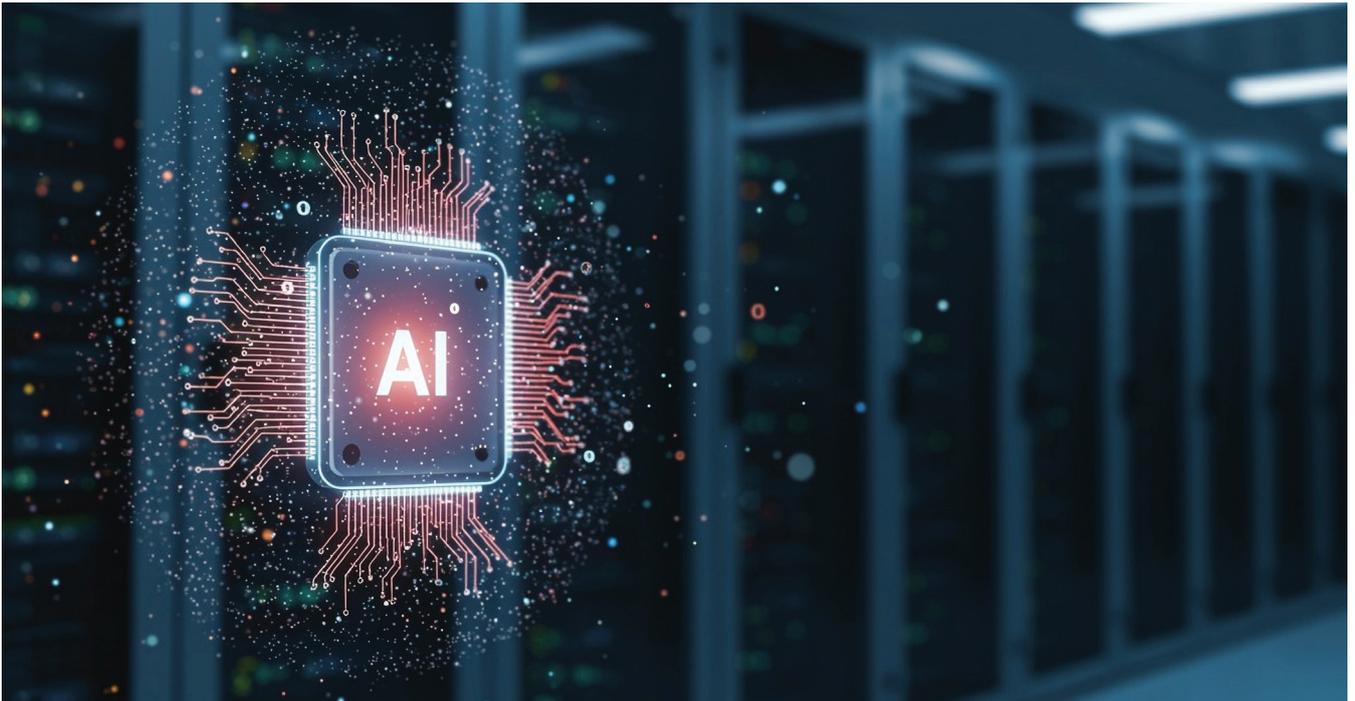
Looking ahead to 2026, the report identifies several trends that will define the sector's next chapter. Demand will continue to place growing pressure on land, water and grid capacity. Water stewardship will become a central concern, alongside a shift from carbon to measuring real ecological outcomes, including habitat and biodiversity restoration.

Community trust will also play a larger role in determining where new data centres can be built. These trends show that nature-led design will need to play a far greater role in how the sector grows.

"The conversation has been framed as a trade-off for too long: innovation versus conservation, progress versus protection. But that narrative no longer holds. Europe's balance between digital progress and nature restoration will rest on the choices of a handful of major technology companies, and it's crucial that they get it right" said Lisett Luik, co-founder of Arbonics.

Vertiv report highlights emerging AI data centre trends

Vertiv's report reveals key technology trends amidst AI advancements, focusing on power efficiency, energy autonomy, and the integration of digital twins.



THE evolution of data centres continues to be influenced by both macro forces and technological trends, particularly those related to artificial intelligence (AI). According to Vertiv's latest report, the Vertiv Frontiers, the industry is evolving to meet the demands imposed by modern-day technological advances.

The report highlights a key change: the shift towards higher voltage DC power architectures to efficiently manage AI-driven workloads. Traditional AC/DC systems are being challenged by increasing power densities, steering the move to more effective power distribution mechanisms.

With a focus on AI data centres, Vertiv foresees these hubs not only handling AI demands but also adapting to the specific regulatory needs of industries like finance and healthcare. As AI becomes central to operations,

businesses face decisions on whether to maintain private environments, especially for secure operations.

Energy autonomy is another key trend. As standalone data centres rely on on-site energy generation for sustainability, widespread challenges in power availability further underline the necessity for self-sufficient energy solutions. Technologies like microgrids and natural gas turbines can support the achievement of this autonomy.

The intersection of AI and data centre design finds its embodiment in digital twins. By using virtual models, data centres can streamline both planning and deployment, accelerating the time-to-market for infrastructure while lowering time arrests by up to 50%. This approach is important in the pursuit of gigawatt-scale expansion.

Meanwhile, AI's role in liquid cooling technologies moves from adoption to refinement of systems. As these cooling solutions are increasingly adopted, AI is harnessed to make these systems even more resilient, predicting malfunctions and ensuring operational continuity.

Vertiv's operates in over 130 countries, bringing solutions that span power, cooling, and IT infrastructure, supporting continuous operation and scalable growth in a digital world. Their integrated array of services seek to support facilities adapt in an ever-complex technological landscape.

The interaction between AI and data centres is expected to influence operational strategies and support the ongoing development of data-driven infrastructure.

EMEA data centre investment on the rise

Data centres in EMEA gain traction amid digital infrastructure expansion, posing challenges and opportunities for investors.

WITH digital infrastructure increasingly central to real estate strategies, data centres are attracting greater attention from EMEA investors, according to Colliers' 2026 Global Investor Outlook.

A surge in capital allocation towards data centres is evident with the industry now encompassing 31% of global capital raised from Q1 to Q3 2025, an increase from an average of 15% since 2020. While the US remains a dominant player, EMEA markets such as Germany, the Netherlands, and the UK are attracting interest due to their favourable connectivity and regulatory environments.

Infrastructure and energy availability have surfaced as decisive factors in selecting sites for new data centre developments. Notably, Frankfurt, Amsterdam, and London show increasing demand, yet face power and

planning challenges that inhibit new projects. To navigate these obstacles, some investors are eyeing alternative locales, particularly in Southern and Eastern Europe, where local governments are investing in digital and energy infrastructure.

The sector is also influenced by strategic platforms and joint ventures. By investing in operational businesses and development initiatives, investors gain expertise and scale, aligning with a broader real estate trend towards active investment strategies and enduring partnerships.

Despite growing demand, several hurdles persist. Reliable energy supply remains a significant concern, particularly in urban areas grappling with grid capacity constraints. Developers also face strict sustainability regulations and intricate planning requirements. Adapting to changing

government policies and innovating design and construction techniques are critical for creating new capacities at scale.

Stakeholders from operators to landlords and investors are called to collaborate, ensuring data centres meet demands for resilience, sustainability, and speed. This synergetic approach helps meet user expectations and capture long-term value in an ever-evolving digital landscape.

As digital transformation hastens, data centre investments will likely remain a priority for EMEA investors. The sector provides growth prospects; however, addressing infrastructure challenges, regulatory intricacies, and operational demands will be key. Investors well-versed in local markets and equipped with strategic alliances stand the best chance of unlocking value in this complex sector.



Data centre programme targets AI-focused startups in Hong Kong

SUNeVision launches its third Startup Programme, fostering innovation in Hong Kong with robust infrastructure and strategic partnerships.

SUNeVision HOLDINGS, data centre provider in Hong Kong and the technology arm of Sun Hung Kai Properties Limited, has announced the third edition of its Startup Programme. By providing digital infrastructure and aligning with technology partners, SUNeVision seeks to establish a comprehensive platform for budding innovators, further solidifying its dedication to enhancing Hong Kong's tech landscape.

Geared towards facilitating growth at every stage, the programme offers technological resources and mentorship to startups. Emphasising the exploration and refinement of AI, with the aim of helping participating startups strengthen their competitive position in sectors such as AI, BioTech, FinTech, GreenTech, HealthTech, Smart City, and more. Open for applications until 10 February 2026, the programme will select 10 promising startups per cohort.

The programme is intended to offer participants a range of benefits, including SUNeVision Credits valued at up to HK\$160,000. These credits are redeemable for various support services from SUNeVision and affiliated partners, laying a strong digital infrastructure to enable AI-driven applications within their data centres. Additionally, startups will access an ecosystem of over 300 technology players and key stakeholders, bolstering their impact in the Greater Bay Area and Asia-Pacific regions. Tailored tools, resources, and mentorship will be dispensed through esteemed programme partners:

- Angelflow: Syndicate technology infrastructure
- Dataplugs: Internet and managed hosting
- Finda Cloud: Value-added cloud services
- Nexusguard: DDoS protection and cybersecurity

SUNeVision supports homegrown talent and accelerates technological innovation, providing an environment for visionary ideas to flourish into tangible impacts on society.

- SUNeVision: Data centre colocation and hosting
- Sustainable SmartTech Ventures: AI-powered smart building management
- The Hive.: Co-working space
- Votee AI: Authentic Cantonese translation
- WeExpand: Agentic AI services



The previous cohort highlighted significant innovators like DeepTranslate and Sustainable SmartTech Ventures (SSV). DeepTranslate melded AI with human insight for financial translations, benefiting from SUNeVision's low-latency network for enhanced accuracy and speed. Meanwhile, SSV utilised AI to advance building management technologies, leveraging SUNeVision's data centres to refine AI models with high-capacity results, thereby attracting enterprise attention in the construction industry.

With such initiatives, SUNeVision demonstrates support for fostering homegrown talent and accelerating technological innovation in Hong Kong, providing a holistic environment for visionary ideas to flourish into tangible impacts on society.

Strategic joint venture develops AI infrastructure in Saudi Arabia

stc group and HUMAIN have formed a joint venture to build AI data centres in Saudi Arabia, supporting high-capacity, low-latency AI workloads and large language models.

stc group, A GLOBAL DIGITAL facilitator, has teamed up with HUMAIN, a company specialising in artificial intelligence backed by the Public Investment Fund (PIF), to launch a Joint Venture (JV). This collaboration aims to construct AI data centres within the Kingdom of Saudi Arabia.

This partnership will support the development and operation of infrastructure designed to accommodate up to 1 gigawatt (GW) of AI workloads. Combining centre3's vast data-centre expertise and regional connectivity with HUMAIN's expertise in AI provides a strong combination for establishing high-capacity, low-latency infrastructure required for the AI revolution.

Olayan Alwetaid, CEO of stc group, asserted that this JV cements their commitment to nurturing the region's digital landscape. HUMAIN's AI capabilities, together with centre3's data-centre experience, provide a foundation for supporting future AI workloads.

Tareq Amin, CEO of HUMAIN, highlighted the JV's focus on developing infrastructure capable of handling high-demand AI workloads and supporting future technological developments.

He emphasised the importance of purpose-built compute systems that provide efficiency and resilience.

Fahad AlHajeri, CEO of centre3, elaborated on their role in actualising the vision. He stressed leveraging their growing data-centre footprint to deliver density, resilience, and availability needed for next-generation AI computing.

Importantly, this JV symbolises an advancement in enhancing infrastructure to accommodate the rising demand in the region.

In essence, this alliance aims to provide high-performance hosting for vital digital assets, bolster the AI ecosystem, and lay a robust groundwork for deploying large language models (LLMs) as well as mission-critical AI applications.

US AI infrastructure set for major expansion

Hut 8 HAS EMBARKED on an ambitious partnership with Anthropic and Fluidstack to accelerate the deployment of hyperscale AI infrastructure across the United States. At the core of this collaboration is the development of state-of-the-art digital infrastructure to propel energy-intensive, next-generation use cases.

The initiative promises to deliver between 245 megawatts (MW) and up to 2,295 MW of AI data centre infrastructure. Using high-performance clusters operated by Fluidstack, Hut 8 aims to meet the growing demands of AI technology.

The partnership unfolds through multiple phases:

- Phase 1: The initial phase will target the River Bend campus in Louisiana. Here, Hut 8 and Fluidstack will establish 245 MW of IT capacity

supported by a 330 MW utility capacity.

- Phase 2: Hut 8 extends the Right of First Offer (ROFO) for up to an additional 1,000 MW of IT potential at River Bend, pending power expansion at the site.
- Phase 3: In a strategic move, Hut 8 and Anthropic will explore the joint development of up to 1,050 MW further, beyond River Bend's existing scope, enhancing the pipeline's capacity.

Asher Genoot, the CEO of Hut 8, emphasised the pivotal role of power in scaling frontier AI infrastructure. He highlighted their innovation-driven model that synchronises power, data centre design, and deployment, creating an integrated, gigawatt-scale platform.

Anthropic's Head of Compute, James

Bradbury, expressed optimism over the expanded collaboration, citing Hut 8's capability to deliver on the scale required, ensuring continuous progress in their model enhancements.

Fluidstack's CEO, Gary Wu, reiterated their commitment to solving compute challenges across various complexities, expressing pride in partnering with institutions like Hut 8 to establish the foundational infrastructure vital for AI advancement.

This collaboration is expected to support the development of AI infrastructure and contribute to growth and technological progress.

Accelerate the deployment of hyperscale AI infrastructure

New sustainable data centres under construction in Spain

Quetta Data Centers is expanding its footprint with the development of eco-efficient edge data centers powered by renewable energy across Spain and Portugal.



QUETTA DATA CENTERS, a project by Azora in collaboration with Core Tech Capital, is set to advance the landscape of edge and AI data centers with its latest ventures. Recently, construction commenced on their Tres Cantos data center in Madrid, doubling its capacity to 20 MW. Simultaneously, their inaugural project in Barcelona is gaining traction in Molins de Rei, both set to strengthen Quetta's presence in the data center arena.

Emphasising sustainability, Quetta is forging ahead with plans to construct six edge data centers across Spain and Portugal. The facilities are designed to incorporate closed-loop liquid cooling, with the goal of reducing water consumption and achieving efficient power usage (PUE) of under 1.15. The projects, backed by over €500 million in investment, aim to advance Quetta's technological capabilities while

maintaining environmental standards. The Barcelona and Madrid expansions are set to elevate Quetta's total capacity to 30 MW. This represents 50% of their overarching target of 60 MW, marking a step on their strategic roadmap. Since being unveiled by Azora in 2024, Quetta Data Centers has merged operational expertise with technology, embodying efficiency, resilience, and sustainability.

The digital economy's meteoric rise has fuelled an ever-growing demand for robust data infrastructure. The sector is rapidly evolving, driven by needs stemming from AI, gaming, augmented reality, and the Internet of Things. In response, Quetta's AI-ready facilities are engineered for high connectivity and low latency, equipped to meet the diverse demands of modern digital workloads. Quetta states that its goals for the data centers are to:

- Operate on 100% renewable energy.
- Ensure no water consumption for an eco-friendly approach.
- Promote uncompromised energy efficiency through innovative solutions.

As Santiago Olivares, a partner at Azora, underscores, there's a significant opportunity in fostering sustainable local data infrastructure, championing assets that fuel economic and environmental advancement. Likewise, Paco González from Core Tech Capital highlights Quetta's commitment to developing a resilient network, offering flexibility and supporting a multitude of clients.

These initiatives position Quetta to support growing demand for AI and digital services while maintaining a focus on resilience and environmental standards.

Accenture strengthens AI data centre services with DLB acquisition

Accenture is set to acquire DLB Associates to enhance its AI data centre capabilities, aiding clients with advanced infrastructure solutions.

IN A STRATEGIC MOVE, Accenture has agreed to acquire a 65% stake in US-based AI data centre engineering and consultancy DLB Associates. This acquisition, involving DLB and its affiliated companies, aims to bolster Accenture's end-to-end data centre capabilities.

Founded in 1980, DLB is recognised for its data centre site selection, due diligence, and design engineering. The firm specialises in commissioning, construction quality management, and energy optimisation services. It has become a partner for some major technology giants and emerging cloud and colocation providers. DLB's innovative approach aims to support clients in rapidly achieving market readiness with successful outcomes.

Once the acquisition finalises, DLB's 620 employees, led by David Quirk and Neil Chauhan, will join Accenture's Industry X practice. This will aim to enhance Accenture's infrastructure and capital projects capabilities. The collaboration hopes to provide clients

with comprehensive solutions from conceptual design through to operational performance.

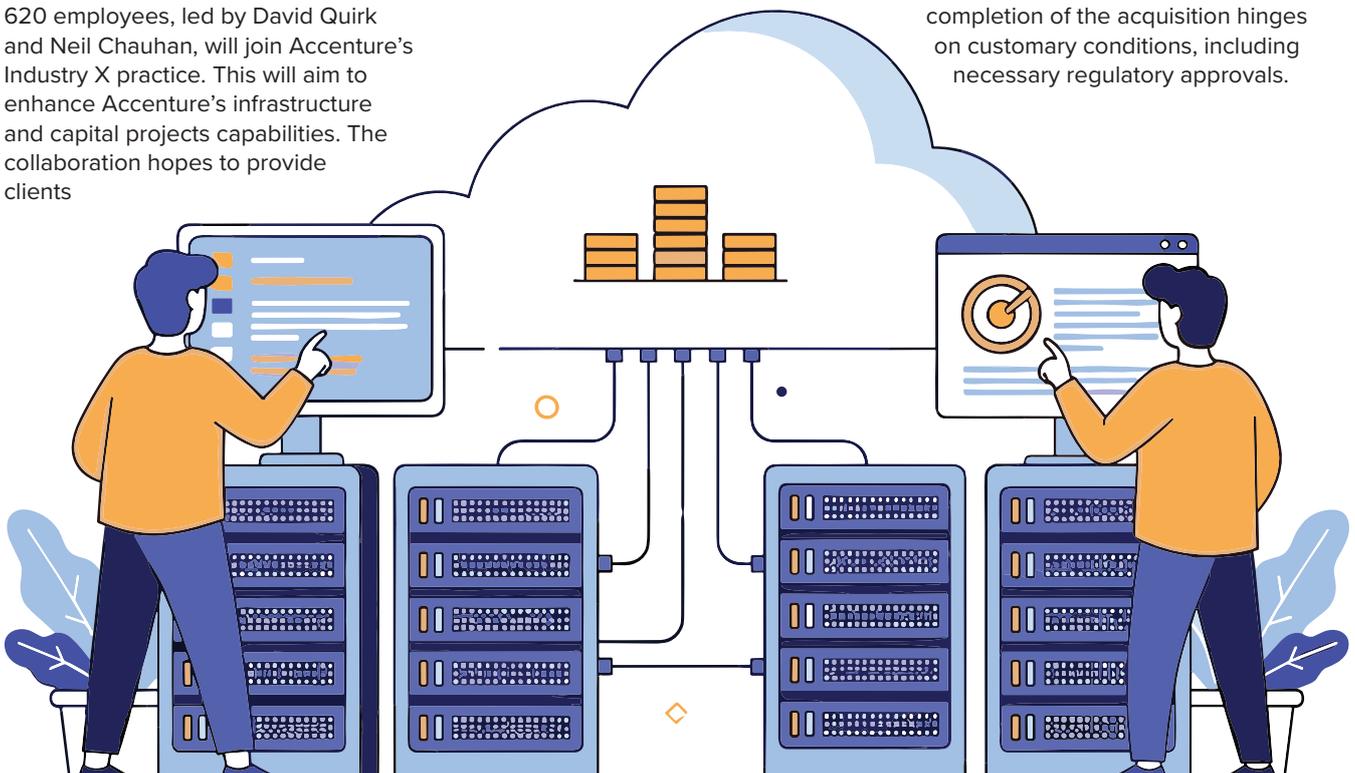
Manish Sharma, chief strategy and services officer at Accenture, emphasised that the merger will help clients access AI-ready infrastructure more quickly and efficiently. He noted that the partnership aims to improve site planning and selection while also addressing challenges such as staffing shortages and cybersecurity risks.

DLB's engagements include pioneering industry standards, evidenced by its co-foundation of the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Technical Committee TC9.9. This committee focuses on the critical aspects of data centres, amplifying their importance.

The data centre industry, fuelled by AI-demand, is on a rapid growth trajectory. David Quirk, CEO, DLB remarked, "We selected Accenture to unleash our ability to match that scale and help support clients through value creation and joint ideation for next generation AI data center—at pace." Confidence in DLB's capabilities to support this scaling endeavour is mirrored by Neil Chauhan, DLB's Chief Innovation Officer.

Accenture's deliberate strides in expanding AI infrastructure capabilities include recent acquisitions of Anser Advisory in the US and Comtech in Canada. These investments aim to enhance Accenture's global standing in infrastructure and capital projects, enabling a diverse service offering.

Though the terms were not disclosed, completion of the acquisition hinges on customary conditions, including necessary regulatory approvals.



Data centre outsourcing market size to cross \$243.3 billion by 2034

As businesses adopt hybrid cloud strategies and edge computing gains traction, data center outsourcing becomes crucial for scalability and compliance.

THE GLOBAL data centre outsourcing market was valued at USD 132.3 billion in 2024 and is projected to grow at a CAGR of 6.4%, reaching USD 243.3 billion by 2034, according to a recent report by Global Market Insights Inc.

The demand for data centre outsourcing continues to rise as businesses increasingly pursue flexible and secure infrastructure solutions. Organizations are embracing hybrid cloud strategies that combine the control of private clouds with the agility of public cloud services. This approach enables companies to scale operations while maintaining tighter oversight of critical data. Outsourcing providers are now offering integrated solutions that span both private and public environments, optimizing performance and cost management simultaneously.

As emerging technologies such as 5G, IoT, and real-time applications gain momentum, enterprises are turning to edge computing for faster processing at the source of data generation. This has led to a shift toward more distributed outsourcing models, where smaller, decentralized facilities are placed closer to end users. In the US, hyperscale operators, including Microsoft Azure, Google Cloud, AWS, and IBM, are leading the outsourcing movement with their massive infrastructure and capacity to support enterprises at scale



without hefty upfront investments. Meanwhile, data privacy frameworks such as HIPAA, FINRA, and CCPA are shaping outsourcing demand, driving businesses to work with providers who offer certified facilities, robust compliance support, and regional regulatory alignment.

The data centre outsourcing market from hardware segment captured 43.7% share in 2024 and is projected to grow at a CAGR of 6.4% through 2034. With rising data volumes and evolving technologies, organizations are opting to outsource hardware management to cut capital expenses and adopt an operating cost model. Managing in-house infrastructure upgrades proves cost-intensive and time-consuming, which is why outsourcing hardware services has become a preferred path for scalability and agility. The power and cooling infrastructure segment is expected to register a CAGR of 8.7% from 2025 to 2034. Outsourcing providers are introducing advanced energy management solutions to support high-performance computing environments. Technologies such as AI-based temperature control, liquid cooling, and free cooling are being adopted to handle heat generated by dense workloads while also reducing energy consumption and enhancing system reliability.

United States data centre outsourcing market held a 76.1% share in 2024, generating USD 34.8 billion. The US remains a global hub for data centres, driven by the presence of major providers such as Equinix, Amazon Web Services (AWS), Verizon Communications, and Google Cloud. As regulatory frameworks become more complex, companies increasingly seek third-party partners with the compliance credentials and infrastructure to navigate

The demand for data centre outsourcing continues to rise as businesses increasingly pursue flexible and secure infrastructure solutions.

evolving data privacy laws. Canada's enterprise market is also transitioning toward cloud-driven outsourcing models, prioritizing speed, innovation, and cross-platform orchestration. Providers with strong hybrid and multi-cloud capabilities are seeing increased traction across the region.

Key companies operating in the global data centre outsourcing market include Cognizant, Tata Consultancy Services (TCS), Fujitsu, Accenture, Amazon Web Services (AWS), Google Cloud, Microsoft Azure, Equinix, Verizon Communications, and Digital Realty. To strengthen their position in the competitive data centre outsourcing space, companies are focusing on expanding global infrastructure, integrating edge computing solutions, and offering hybrid and multi-cloud management platforms. Strategic investments are being made in AI-based automation for data management, energy optimization, and real-time monitoring. Providers are also forming alliances with hyperscale cloud vendors to co-deliver scalable services while ensuring compliance with evolving regional regulations. Emphasis is being placed on offering flexible service models, cost-effective infrastructure-as-a-service (IaaS), and dedicated support for industry-specific compliance, like healthcare or finance.

Innovative solar project marks new era for UK data centres

A pioneering solar farm collaboration aims to transform data centre energy consumption through renewable solutions.



DOWNING Renewable Developments (DRD) has embarked on a groundbreaking venture within the data centre sector, launching its inaugural initiative dedicated to sustainable energy solutions. A significant agreement has been reached with Kao Data to develop, build, and operate a pioneering 40MW solar photovoltaic (PV) farm. This Green Data Solar Farm project aims to supply Kao Data's renowned data centre campus in Harlow with up to 40 megawatts of solar-generated electricity.

One of the notable aspects of this initiative is its role as a potential catalyst in Kao Data's commitment to reducing carbon emissions and enhancing energy resilience. By integrating verified renewable sources, Kao Data takes a definitive step towards becoming fully net zero by 2030.

The project's significance is further enhanced as it becomes one of the first ground-mounted solar PV systems to directly power a UK colocation

data centre campus. Through a Power Purchase Agreement (PPA), Kao Data is securing long-term access to this renewable energy source, marking an important partnership aimed at further decarbonising data centre operations.

This collaboration comes at a crucial moment as the UK addresses the dual challenges of systemic energy demands and soaring energy costs. Proactive measures from entities like Kao Data demonstrate how the private sector can offer innovative solutions for energy availability and security.

Renewable energy partnerships between DRD and Kao Data spotlight the ability to meet technological and economic ambitions outlined in the Government's AI Opportunities Action Plan. The Green Data Solar Farm is an exemplary demonstration of how private wire renewable energy systems can propel the UK's green transition. Such initiatives encourage data centre operators to adopt a prosumer model, paving the way for self-generation of

renewable energy, securing long-term low power costs, and reducing grid demand.

As described by Tony Gannon, Downing Renewable Developments sees this project as a milestone. It underscores the growing importance of integrated clean energy solutions in digital infrastructure, setting new standards for meeting the energy needs of mission-critical assets such as data centres.

Mathew Harris of Kao Data adds that decisive action and innovation are pivotal to maintaining the UK's competitive position in the AI and digital sectors. The strategic collaboration with DRD provides a blueprint for addressing renewable energy challenges, delivering green power rapidly and at scale to fulfil immediate requirements.

This pioneering project showcases how the collaboration of private sector entities and local authorities can be instrumental in advancing sustainable energy and economic progress.

Rethinking AI Infrastructure: Why green data centers must be chosen responsibly



AI and data centres are not neutral tools - they are embedded in systems that affect water, energy, climate and equity. We cannot afford to leave those systems unexamined or unaccountable.

BY JEAN-MARC BOURREAU, CO-FOUNDER, AFRIK FOUNDATION

THE UK Environment Agency recently issued a warning: the rapid growth of artificial intelligence is making water scarcity harder to predict (1). The rise in digital processing power, particularly from data centers supporting AI models, is putting unpredictable strain on public water supplies. In the UK, where many data centers rely on municipal water for cooling, this is happening without proper oversight or measurement.

This is not just a British concern, it's a global signal that the infrastructure powering our digital future is being built without enough transparency or sustainability in mind. If left unchecked, digital expansion risks accelerating the very environmental crises it's supposed to help solve.

At Afrik, we do not build or operate data centers. Instead, we focus on the democratic governance of infrastructure. Our platform empowers communities to vote on which projects – such as renewable energy, climate-adaptive agriculture or sustainable digital infrastructure – should receive funding. Using blockchain technology and smart contracts, we provide tools for transparent, community-led investment that is geo-weighted and locally accountable.

AI models and digital tools can be used to fight climate change – but only if the infrastructure supporting them is equally responsible. When data centers are built without regard for energy sources or water use, their

environmental toll can be severe. Yet too often, local communities have little to no say in where or how these centers are built.

Afrik was created to address this imbalance. Through our voting platform, communities and stakeholders can make informed decisions about proposed infrastructure projects, ensuring that digital development aligns with local environmental priorities. In regions considering green data centers or energy-intensive digital processing hubs, we see our role as a facilitator of accountable funding and public oversight, not an infrastructure operator.

This model allows communities to prioritize low-carbon energy sources, monitor environmental metrics and ensure any AI or blockchain applications are locally beneficial and not extractive.

Our platform supports smart contracts – programmable agreements that unlock funding only when agreed-upon milestones or ethical criteria are met. For example, a community could vote to fund a renewable energy microgrid, contingent on using local labor or meeting biodiversity standards. Or a green data center could be funded if it meets specific energy or water efficiency benchmarks.

We call this geo-weighted governance, it's a model in which those most affected by development decisions have the greatest say and is a sharp





contrast to top-down infrastructure planning that often leaves out the people it impacts most.

One of the clearest use cases for community-governed digital infrastructure is agriculture. AI can help smallholder farmers adapt to climate shifts, optimize planting cycles and access early warnings for droughts or pests.

Afrik's platform enables funding for farmer-focused digital tools, many of which use AI to support productivity and sustainability.

We believe food sovereignty requires data sovereignty. Communities must be at the center of decisions about how AI is used on their land, and whether digital infrastructure enhances or erodes their livelihoods.

Every decision about digital infrastructure is also a climate decision. The placement of a data center, the source of its electricity, the amount of water it consumes, all have long-term environmental consequences.

That's why we advocate for a shift in mindset: digital infrastructure should be viewed as climate infrastructure.

Through blockchain-based voting and smart contract enforcement, communities can build mechanisms that align digital expansion with sustainability goals. This includes public visibility into environmental impacts, inclusive governance frameworks and long-term accountability for outcomes.

At Afrik, we also address financial exclusion through the creation of digital assets that reflect local value, not global speculation. In many regions, weak or volatile currencies make it difficult for communities to fund essential infrastructure. Our approach ties digital tokens to real-world, productive activity, thus enabling transparent, direct participation in solar farms or digital infrastructure.

We believe poverty is not just a lack of capital, but a lack of control over capital flows. By decentralizing funding decisions and linking value to impact, we help restore that control to local communities.

Our goal is not to impose technology, but to steward it. That means putting in place ethical guardrails, transparent governance mechanisms, and metrics that track both social and environmental performance. From rural Africa to urban

Europe, the same principle holds: people must have agency over the systems that shape their lives.

The warning from the UK Environment Agency is a wake-up call. AI and data centers are not neutral tools - they are embedded in systems that affect water, energy, climate and equity. We cannot afford to leave those systems unexamined or unaccountable.

Afrik is helping to build the digital governance tools we'll need for a more inclusive, sustainable future. These tools allow communities to say yes to the right projects, and no to the wrong ones. But this cannot be done single handedly. Governments, businesses, technologists and civil society must work together to ensure the next wave of infrastructure serves people and the planet alike.

It's time to slow down, ask the right questions and make decisions in public. Because digital infrastructure isn't just shaping our future, it's shaping our climate, our economies and our democracies.

<https://www.globalgovernmentforum.com/ai-makes-water-shortages-harder-to-predict-says-uk-environment-agency>

We believe food sovereignty requires data sovereignty. Communities must be at the center of decisions about how AI is used on their land, and whether digital infrastructure enhances or erodes their livelihoods



Can DCIM’s AI-driven future increase data centre sustainability?



The rapid expansion of AI, big data, and IoT has intensified the complexity of data center capacity and infrastructure management. However, with data center infrastructure management (DCIM) solutions’ digital intelligence, operators can integrate the latest advanced technologies and simplify their mission-critical facility operations.

BY DAVE SCHAIBLE, COO OF RADIX IOT

THIS HAS BOOSTED the DCIM market, which is expected to reach \$5.01 billion by 2029. By 2030, US-based companies are expected to invest over \$1.6 trillion in capital expenditures on hyperscale data centers and AI computing infrastructure, representing over 40 percent of the \$7 trillion projected global total.

DCIM keeps data centers’ energy appetite in check

Considering that in 2024, the global data center energy usage accounted for nearly 1% of total energy consumption, U.S. data centers’ energy consumption is projected to reach up to 12 percent by 2028.

With large AI models consuming energy equivalent to thousands of homes, data center operators now face heightened scrutiny over the energy consumption surge, environmental impact, carbon footprint, and governance.

By integrating DCIM solutions, operators can efficiently manage and keep energy consumption in check with real-time visibility into critical data across multiple devices and systems.

As substantial volumes of data are analyzed in real-time, operators can make proactive, data-driven business decisions to enhance equipment performance, and preempt failures for operational efficiency and environmental sustainability.

This has transformed DCIM solutions from merely a reporting tool into a proactive digital intelligence network that helps:

- Optimize resource allocation and energy consumption efficiency by analyzing workload patterns and adjusting infrastructure accordingly
- Proactively control energy consumption, respond to alerts

signaling power usage discrepancies, and reduce energy-intensive maintenance resulting from equipment failures

- Optimize power and cooling needs and preempt potential failures by flagging proactive maintenance and planning

In identifying energy inefficiencies, operators can reduce downtime, improve operational efficiency, increase sustainability, and ultimately enhance data center resilience in managing fluctuating loads. This allows operators to:

- Set alerts and adjust settings to collect necessary data for on-demand validation
- Tack trends over time to improve decision making and reliability
- Establish energy and sustainability goals and adjust the system in real time to avoid future overspending

Operators can also use DCIM solutions to set and achieve long-term strategic goals, while maximizing operational impact, by:

- Minimizing energy consumption and meeting Service Level Agreement (SLA) compliance
- Streamlining ESG reporting by collecting comprehensive data sets on sustainability metrics, including carbon emissions, energy consumption, water use, and fuel use for generators
- Building dedicated ESG dashboards for on-demand data retrieval and automated report generation.

As more data centers host AI applications, DCIM solutions become more critical tools for monitoring power demands and improving data feed latency. Considering how small increases in operational load spike megawatts of grid power, scalable DCIM software's real-time telemetry data can provide full visibility across all facilities' systems, so operators can shift workloads, balance power loads, eliminate power grid overloads or risk costly outages.

In expediting data collection from a range of systems and across globally distributed servers, DCIM solutions also enable operators to increase data collection speeds without any bottlenecks.

Optimizing physical resources

With the proliferation of multi-site data center operations, DCIM solutions are increasingly essential for addressing the challenges of remote monitoring and provisioning in decentralized environments, particularly for edge computing. Operators can efficiently monitor the operational performance and track assets in real-time from dispersed edge infrastructure. They can also streamline workflows and proactively address risks to equipment, security, power, and cooling systems to ensure greater uptime while avoiding costly downtime.

DCIM solutions also allow operators to optimize physical resources across

data centers as their digital workloads spike. This allows for effective planning for allocated space, power, and cooling needs, and new equipment requirements, while preventing outages and costly downtime.

With automated tasks and alerts, operators can be notified and proactively resolve flagged issues remotely. This expedites efficient and cost-effective management of distributed sites and assets by using a single, holistic platform. DCIM's real-time alerts facilitate access to critical datapoints and enable timely responses and repairs, enhancing overall risk mitigation with remote monitoring and troubleshooting.

Safeguarding against elevated cyberthreats

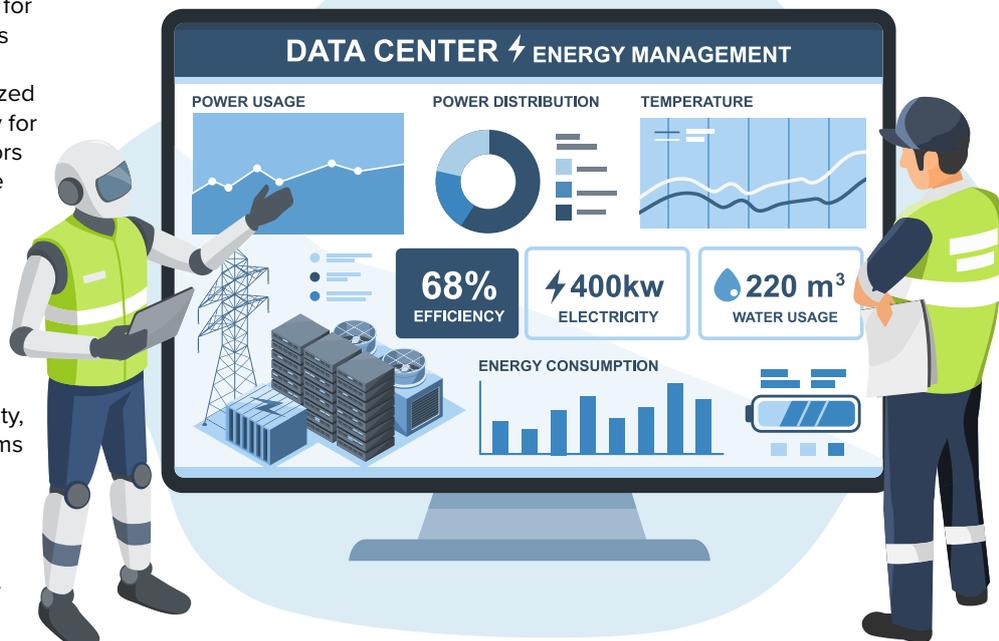
As most data centers adopt hybrid cloud environments and utilize third-party vendors, the risk of cyberattacks such as data theft, malware, phishing, and data leaks has significantly increased. A global study by Statista found that data breaches cost on average \$4.88 million per incident between May 2020 and February 2024.

With 100 percent uptime non-negotiable for all data centers that are enabling AI's massive-scale computations, a cyberattack targeting any aspect of a data center's operational

As more data centers host AI applications, DCIM solutions become more critical tools for monitoring power demands and improving data feed latency. Scalable DCIM software's real-time telemetry data can provide full visibility across all facilities' systems

technology (responsible for all equipment controls of cooling or power) is a costly risk that can result in instant failures of components critical to AI operations.

In addition to managing security, ensuring compliance, and enhancing threat detection, DCIM solutions are operators' security and defense shields against cyber threats. In strengthening the infrastructure through physical access monitoring to detect hardware behavior anomalies, DCIM solutions help automate threat detection and safeguard data center infrastructure from cyber-attacks.



More importantly, DCIM solutions must be integrated with the facilities' physical security. To ensure systems are up to date, operators must implement strong access controls and encrypt communications. Segmenting networks can also further reduce attack surfaces.

DCIM Solutions future-proof data center infrastructure

DCIM solutions are a data-driven intelligence system that prompts tactical business decision-making. Operators rely on DCIM solutions to forecast capacity by collecting and analyzing historical and real-time data on their space, power, cooling, and compute resources. By analyzing this broad data set, operators can identify trends and forecast future needs more effectively.

In eliminating antiquated, manual spreadsheets, DCIM solutions' real-time, holistic view of a data center allows for predictive CapEx forecasting and lifecycle planning by aligning IT and facilities teams. It also provides ROI

calculations that can be readily turned into segmented reports to meet the specific needs of various stakeholders.

In our hyperconnected, 24/7 digital reality, made possible by data centers' optimized uptime, downtime risks are not an option when, on average, downtimes cost an estimated \$9,000 per minute, and can easily exceed \$1 million per hour in the case of major enterprises.

Without DCIM solutions, organizations might as well be functioning in the dark, lacking full visibility and control over their resources. Not to mention that being in a reactive, costly, and less agile approach to facilities management is not an option for a profitable business.

The consequences of increased downtime, operational inefficiencies, and higher costs due to under-or overprovisioning resources will eventually lead to capacity planning

On average, downtimes cost an estimated \$9,000 per minute

challenges that can further hinder sustainability compliance.

In the fragmented data center infrastructure environment, where operators face daily challenges and risks due to siloed data, DCIM solutions remain essential intelligence tools that reliably scale to unify data from power and cooling systems.

Ultimately, in providing visibility into real-time data analytics of all operational systems, DCIM solutions enable proactive alerts that allow operators to reduce downtime, lower energy consumption, and significantly impact the bottom line without losing sleep about the unpredictability of "what-if" scenarios.



DCS **DATA CENTRE SOLUTIONS ROADSHOW**

30 Seats. 1 Day. VIP DC Strategy Discussion
LIMITED PLACES AVAILABLE – APPLY TODAY!
24 March 2026

THE DCS ROADSHOW 2026 is an exclusive executive forum, limited to 30 senior leaders responsible for data centre ownership, development, power strategy, and delivery across England.

The audience includes operators, hyperscale infrastructure teams, utilities, developers, EPCs, OEMs, infrastructure investors, and a limited number of independent advisors.

Apply for your complimentary pass to the **DCS Roadshow 2026** in London at: <https://datacentreroadshow.com/events/london-2026/register>

Interested in speaking or attending? If you have any questions about attending as a delegate or speaker, please reach out to info@datacentreroadshow.com

WHAT MAKES THE ROADSHOW UNIQUE?

- **Peer-to-peer Learning:** First-hand insights from industry leaders
- **No Vendor Sales Pitches:** Strategy-led discussions only
- **Curated Networking:** Build meaningful, high-value connections
- **Intimate Format:** Just 30 delegates for focused collaboration



CELEBRATING 16 YEARS OF EXCELLENCE IN THE DATA CENTRE INDUSTRY

Now in its 16th year, the DCS Awards return to honour the innovation, leadership, and achievements shaping the future of the data centre sector.

Reflecting a Rapidly Evolving Industry

Today's data centres are transforming fast, driven by sustainability goals, AI-powered optimisation, and the rise of edge and hybrid architectures. Operators are embracing renewable energy, improving efficiency, strengthening security, and adopting zero-trust and AI-driven threat detection. ESG initiatives feature prominently as the industry continues to underpin the digital world.

Recognising Innovation and Leadership

With more than 300 nominations each year, the DCS Awards celebrate standout projects, technologies, services, and individuals across 38 categories spanning both facilities and IT. Winners are revealed at a prestigious evening event in London, bringing together industry leaders for networking, entertainment, and fine dining.

**NOMINATIONS ARE OPEN
& FREE TO ENTER**

KEY DATES:

6 OCTOBER '25	Nominations Opened
6 FEBRUARY '26	Nominations Close
13 MARCH '26	Shortlist announcement
16 MARCH '26	Voting Open
17 APRIL '26	Voting Close
14 MAY '26	Ceremony

Winners will be announced at a gala evening on 14 May 2026 at Leonardo Royal Hotel London St Pauls, London



Get Involved

Nominations are open for organisations, teams, partners, and customers. Sponsorship and promotional opportunities are also available, offering extended visibility for early participants. Save the Date

NOMINATE NOW!

<https://dcsawards.com/nominate>



Partner data centres with biodiversity for a sustainable AI era



Do we want to be remembered as an industry that opposed our natural world and caused extinctions, or one that set an example of industry working alongside nature?

BY DAVID DAVIES, ASSOCIATE DIRECTOR AT ARUP

DESPITE THE FEAR of a potential AI 'bubble' bursting, data centre development is showing no sign of slowing down in the UK. Recent predictions put data centre spending at over £10 billion by 2029, while almost 100 new facilities are currently progressing through the UK's planning system. The government is committed to expanding its AI capabilities and becoming an 'AI superpower', and if it aims to keep up with global demand and technological advancement, the UK has no choice but to continue investment in data centres.

There is no shortage of discussion around the energy demand and carbon emissions challenges associated with the growth of data centres. The new UK Net Zero Buildings standard and the EU's updated Energy Efficiency Directive reflect a shift towards

overcoming these concerns. However, often left out of the conversation is the effect data centres have on nature.

We are living in a global biodiversity crisis, with research showing humans are driving biodiversity loss among all species. While this topic tends not to be at the forefront of UK sustainability conversations, our ecosystems are at serious risk from rushed infrastructure development. Many data centres fit into this category, often being built on large greenfield sites and risking the loss of crucial habitats.

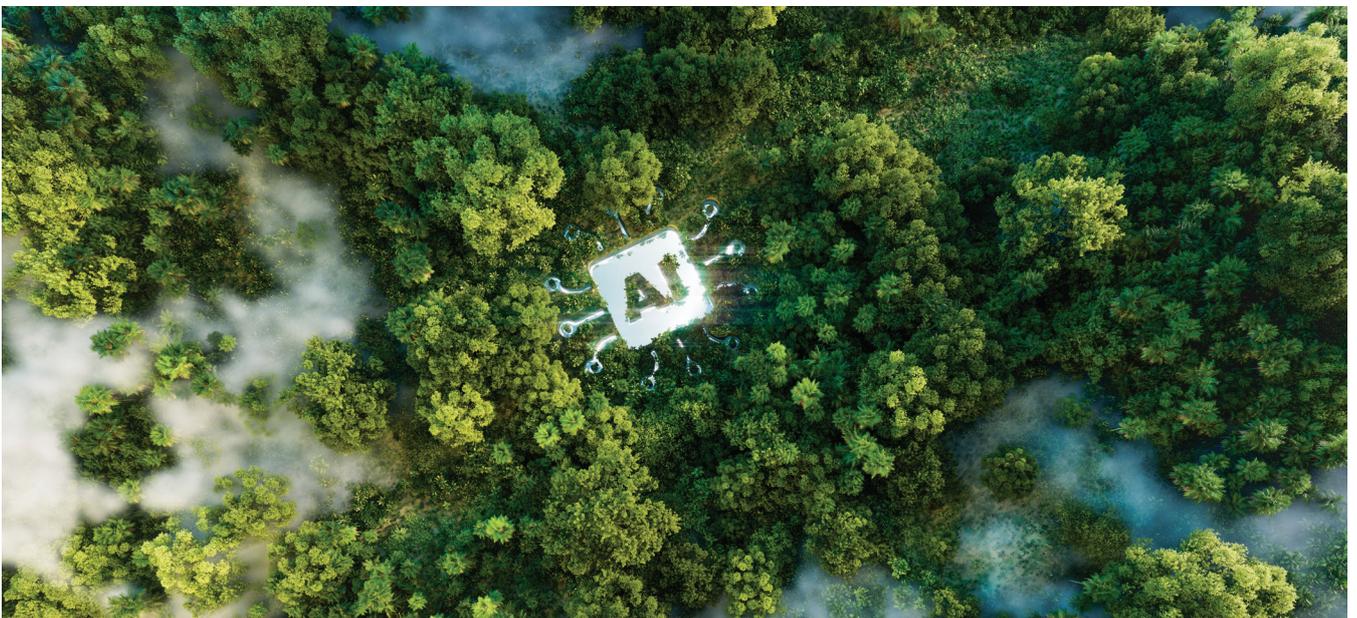
If the UK is to ride the AI wave sustainably and responsibly, we must limit the damage of these developments. Luckily, there is reason to be optimistic. When delivered responsibly, data centres do not have to damage their environments. In fact, they

can actually reverse nature damage and boost biodiversity.

Evolving regulation – the first incentive

Data centres may soon be forced to consider their biodiversity impact due to the changing regulatory landscape. While the UK charts its own course post-Brexit, the EU's Nature Restoration Law – mandating biodiversity recovery by 2050 – signals the direction of travel for environmental regulation globally. Similar frameworks are being implemented by different authorities across the globe, and the UK won't want to lag behind in its environmental ambition.

In practice, this means data centre operators must move beyond traditional environmental impact assessments and toward demonstrating genuine





ZincFive

BC 2 AI

**Designed for AI.
No Tradeoffs in
Safety, Space,
or Sustainability.**

LEARN MORE

zincfive.com

ecological contribution. Working collaboratively with ecologists can ensure data centres actively support and promote their surrounding natural environment.

The business opportunity – the second incentive

Many developers are also recognising that ecological enhancement isn't simply environmentally responsible – it's commercially advantageous. Biodiversity-positive projects face fewer planning delays, generate less community opposition and attract tenants increasingly focused on their environmental, social and governance (ESG) credentials.

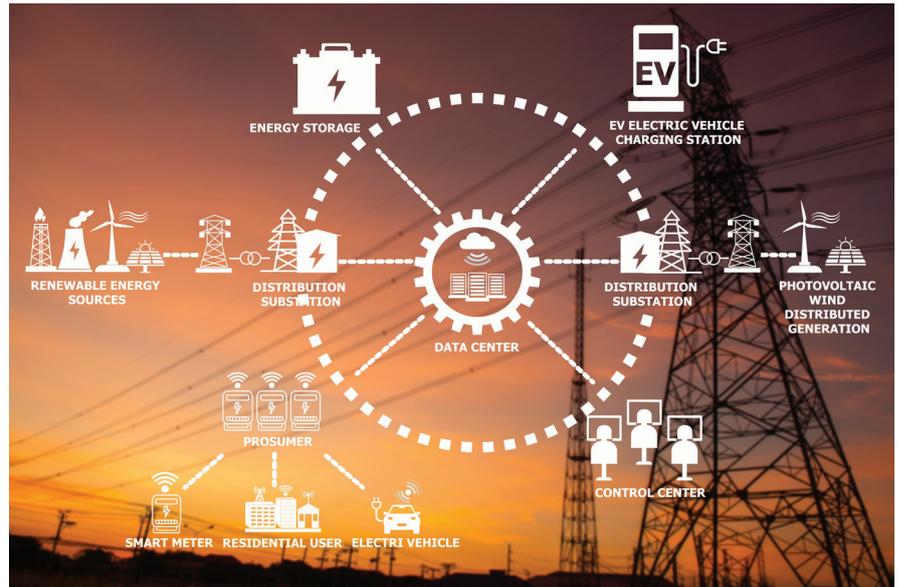
This approach also builds resilience. Climate change brings increasing risks from extreme weather events, water scarcity and temperature fluctuations. Data centres designed with natural systems integration will be better positioned to adapt to these challenges while maintaining operational reliability. With planning authorities increasingly scrutinising biodiversity outcomes, and communities demanding more in their environmental expectations, forward-thinking developers recognise this isn't just about compliance. It's about future-proofing investments and building genuine social licence to operate in an era of heightened environmental awareness.

What can data centres do?

The data centre industry is already well-versed in systematic environmental assessment through carbon accounting. The same rigorous, evidence-based approach can - and should - be applied to nature.

By adopting frameworks like the Taskforce on Nature-related Financial Disclosures (TNFD), businesses can incorporate nature into their financial and business decisions. At Arup, we have been using similar methodologies for years to help projects evaluate and reduce carbon emissions across the value chain. Now, we're applying that same thinking to nature.

This methodological shift is already yielding results. Assessment approaches that evaluate nature-related risks, dependencies and opportunities across the entire data centre value chain are revealing significant potential



for ecological enhancement that traditional environmental impact assessments often miss.

The key is moving beyond checkbox compliance toward strategic integration of ecological considerations into core business decisions – from site selection through to operational management. But what do the operational solutions look like in practice?

The opportunity for data centres to benefit their environments is more extensive than many realise: sustainable drainage systems can recreate wetland habitats while managing surface water; green roofs and living walls provide nesting sites and food sources for urban wildlife; and native tree planting creates visual screening while establishing wildlife corridors.

A recent Arup project in Spain illustrates the potential of this approach. Working across multiple disciplines – including ecology, landscape design and water management – the development team created environmental interventions that recovered local species, improved habitat connectivity for animal movement and restored degraded watercourses. The result was measurable biodiversity net gain and a facility that became a genuine community asset rather than an unwelcome neighbour.

These approaches aren't confined to rural locations. Urban data centres can reduce local temperature extremes through strategic landscaping, while utilising plants and their natural ability to

filter and purify air to improve air quality. Finally, the added benefit of extra green spaces benefits both wildlife and local communities.

An AI superpower without costing the planet

The UK has ambitions to become a leader in the AI space. While this is commendable, it is not enough on its own. It must coincide with becoming a pioneer in promoting data centres working with nature, not against it. In so many industries, we now understand that progress cannot come at the expense of the planet, and data centres are no exception.

The crisis the world's ecosystems are facing requires us to take on this challenge urgently. If those of us working in the data centre industry want to be remembered for a boom that brought about a positive change to the world, we must look to partner data centres with nature. This means going beyond abiding by current laws, but ensuring teams across disciplines are working to identify opportunities to help restore nature, not just limit damage. This must be done early on in project design and as a central element, not as a side note. Operators that contribute to biodiversity net gain will stay ahead of their competition, and reap the commercial benefits – however, for me, the biggest incentive is the legacy we leave behind. Do we want to be remembered as an industry that opposed our natural world and caused extinctions, or one that set an example of industry working alongside nature?



New blog from AFL

Hyperscale Market Shifts: AI, Neoclouds, and the New Limits of Data Centers

AI is driving a new era of hyperscale data centers built for continuous, large-scale workloads. Rack-scale GPUs and high-speed fabrics (400G, 800G, and beyond) are accelerating changes in cluster size, topology, and fiber density.

Read AFL's blog: [Hyperscale Market Shifts: AI, Neoclouds, and the New Limits of Data Centers](#). See how AI is influencing hyperscale infrastructure, and what these shifts mean for fiber, network design, and physical planning.

Read the blog now:

[Market Shifts: AI, Neoclouds, and the New Limits of Data Centers](#)

<https://www.aflhyperscale.com/articles/hyperscale-market-shifts-ai-neoclouds-and-the-new-limits-of-data-centers/>



What's coming down the utility track in 2026/27...



The following is a brief summary of the main points raised by Richard Koszykowski, an energy advisor to the MEUC at the excellent 10x10 utility presentations seminar held by the DCA last month in London. Covering key future issues and event in 10 minutes is a challenge but it focusses the mind on what really is important for energy buyers to focus on in the coming years.

BY RICHARD KOSZYKOWSKI, ENERGY ADVISOR – MEUC

ALTHOUGH the degree of price volatility has continued to decline over the past year, volatility will continue to remain a key issue for 2026. The key drivers on energy prices will continue to be.

- National demand for energy
- Geopolitical issues, such as Russia/ Ukraine, the Middle East, and now Venezuela and Greenland Weather
- Availability of gas from Norway to the UK, but also Russia's oil and LNG to India and China
- Donald Trump

Without doubt, the somewhat wild and erratic proclamations and acts by the US president have had some of the biggest impacts on the price of gas and power throughout the world in 2025, and will continue to do so in 2026.

Today, prices for gas and electricity 1 and 2 years out are lower than the market price for this year. And unfortunately, that is the only good news.

The main issue over the coming years will be the rising level of non-energy costs the majority of commercial and public sector organisations are going to pay. Those companies fortunate enough to be eligible for the Energy Intensive Industries (EII) scheme will be able to get even greater exemptions from an already significant range of non-energy levies. The government also intends to provide similar

exemptions to several thousand manufacturing companies included in the forthcoming BICs scheme, designed to significantly lower electricity costs for energy-intensive manufacturing industries, such as chemicals, automotive, and aerospace.

The cost of these exemptions will be paid for by additional non-energy levies placed on all other commercial and public sector organisations energy bills. And, along with the annual increase in the Climate Change Levy (CCL) Rates

the following new levies are being introduced.

- Nuclear Regulated Asset Base (Sizewell C)
- Hydrogen Production Levy

Furthermore, the cost of transmitting (TNUoS), distributing (DUoS) and using electricity to and on your sites will also be rising dramatically over the coming 5 years to support the government's electrification and clean energy strategies. In London, over the next 2 years there will be little to no change in DNO fixed or time of use charges.

However, site capacity charges will typically increase by well over 100%..... All DNO prices are published 15 months ahead of introduction and available on the web.

Over the next 5 years, annual fixed TNUoS transmission charges, available on the web are forecast to rise for;

- Low voltage (1) sites from £1500 to £4500 pa
- High voltage (1) sites from £8000 to £24,500 pa
- Extra High Voltage (4) sites from £1.7m to £4.1m pa.

It is forecast that within the next few years the ratio between energy wholesale and non-energy costs will change for 50/50 to 30/70. An effective procurement strategy will remain key, but managing non-energy costs will become increasingly challenging.

The cost of transmitting (TNUoS), distributing (DUoS) and using electricity to and on your sites will also be rising dramatically over the coming 5 years to support the government's electrification and clean energy strategies



A practical guide to green energy procurement for data centres



As pressure builds for data centres to decarbonise, green energy procurement has become a strategic priority. If you're exploring ways to cut the carbon footprint of your facilities, manage rising power costs, or navigate the complexities of renewable energy sourcing at scale, our quick guide will help you get started.

BY BEN BEETHAM, SUSTAINABLY SPEAKING HOST, SUSTAINABLE ENERGY FIRST

GREEN ENERGY PROCUREMENT (a term used interchangeably with 'renewable energy procurement') is when a business or organisation purchases energy from renewable sources to meet their energy needs.

This reduces their reliance on fossil fuels, and subsequently lowers their carbon impact. Green energy sources typically include wind, solar, and hydro.

There's a strong business case for switching to green energy, including:

- Compliance with growing sustainability regulations
- Enhanced ESG credentials
- Meeting stakeholder expectations and investor demands
- Price certainty
- Business resilience
- Competitive advantage

While direct energy consumption is classed as a Scope 2 greenhouse gas (GHG) emission, businesses with ambitious net zero targets are now looking beyond the emissions within their direct control and are aiming to decarbonise their entire value chain (Scope 3 emissions).

As such, supply chain partners are under increasing pressure to set net zero targets of their own – sourcing renewable energy is one way of doing this.



The renewable energy procurement hierarchy

However, not all renewable energy is equally as 'green', with key impact differences in the source of energy, the technologies used and the mechanisms used in procurement.

The renewable energy procurement hierarchy always prioritises sourcing energy from the most impactful and greenest sources first.

- **On-site generation**
On-site technologies such as solar

panels and wind turbines on business premises.

- **Corporate purchase power agreements (CPPAs)**

With "indirect" or "virtual" CPPAs, energy generation can be located on or off-site. Users can negotiate directly with the renewable developer to purchase power at a fixed price over a long period of time – usually 15-20 years.

- **Sustainable energy first's sustainable energy consortium**
A straightforward way for businesses

to access traceable, certified green energy from multiple renewable sources under a five-year fixed-price contract.

REGO-backed tariffs

Renewable Energy Guarantees of Origin (REGOs) are certificates that prove electricity has been generated from renewable sources. These have variable impact, as some suppliers will buy directly from renewable sites or build their own, while others will buy standard fossil fuel energy from the grid and pair it with REGO certificates (known as unbundled REGOs) which can be purchased or sold separately to the renewably-generated power.

Understanding CPPAs for business

For businesses looking to meaningfully green their energy supply but are unable to install onsite renewables, Corporate Power Purchase Agreements are often the next most effective route. They provide long-term price certainty, support the development of new renewable generation, and deliver traceable, high-quality green energy.

That said, CPPAs can be complex and aren't always a fit for every business. They typically involve long-term contracts and may require a certain level of demand, creditworthiness or in-house expertise – factors that

The SEC offers a more flexible, low-barrier way for businesses to procure traceable, high-quality green energy from a mix of renewable assets, including solar, wind and hydropower

make them more suitable for larger organisations.

For businesses that want to make a real impact but need a more flexible alternative, there's another option: Sustainable Energy First's Sustainable Energy Consortium (SEC).

The Sustainable Energy Consortium as an alternative to CPPAs

The SEC offers a more flexible, low-barrier way for businesses to procure traceable, high-quality green energy from a mix of renewable assets, including solar, wind and hydropower. Designed to democratise access to PPAs for UK businesses with electricity consumption over 1GWh per year, SEC enables businesses to join a competitively-priced contracting model under a fixed price contract of just five years and access traceable deep green energy without the complex, time consuming and costly contracting hassle that can come with CPPAs.

As well as a short-term contract, benefits of SEC include direct energy sourcing, pass-through non-commodity charges (truly cost reflective rather than fixed with a premium), and half-hourly matching, where energy used is matched with generation from specific assets, making that matched supply 100% traceable.

Choosing SEC also means contributing to investment in new renewable energy generation, and supporting local community projects and UK economic development. Neither a CPPA nor a standard green energy contract, SEC offers the best aspects of both, enabling businesses to make a tangible and verifiable difference to their energy impacts in a straightforward and accessible way.

THE DCA (DATA CENTRE ALLIANCE) IS THE UK TRADE ASSOCIATION FOR THE DATA CENTRE SECTOR

WE SUPPORT DATA CENTRE OWNERS & OPERATORS, SUPPLIERS ALONG WITH ALL THOSE WORKING IN THE SECTOR



We ensure the DC community stays informed and connected through our events, roundtables, articles and more.

Contact The DCA to find out more about the benefits of our Partnerships

t: +44 (0)845 873 4587 · e: info@dcauk.org
www.dcauk.org



Powering the AI revolution – Wesco Anixter’s strategy for accelerating data centre deployments



At the PowerEx Live DCA 10x10 Briefing event held on 11 December 2025 in London, Craig Doyle, EMEA Senior Director of Sales for Wesco Anixter Data Centre Solutions, presented a compelling overview of the rapidly evolving power and infrastructure challenges facing data centres in the age of AI.

WITH DECADES of experience across data centre and telecommunications ecosystems, Doyle highlighted how surging demand for compute capacity is reshaping the way developers plan, build, and scale missioncritical facilities.

One of the most pressing issues he identified is the global scarcity of power availability. Data centre operators increasingly secure land and initiate construction only to encounter significant delays in securing grid connections.

This misalignment between development speed and utility readiness is driving innovative alternatives, such as natural gaspowered generation to bridge the supply gap. In other cases, companies are strategically moving new builds to regions with more abundant available power.

These shifts place a renewed emphasis on robust mediumvoltage systems – whether owned by utilities or by the data centre operators themselves – to ensure resiliency, reliability, and future scalability.

As facilities now routinely scale beyond 100 MW, he highlighted the role of streamlined procurement and strategic inventory planning in supporting aggressive deployment schedules. Speed and safety were also recurring themes throughout the presentation.

Doyle pointed to innovations such as enhanced cable support systems and reusable cable ties that are engineered to handle delicate fibre infrastructure while accelerating installation.

These technologies, paired with expert logistics coordination, help contractors avoid delays, reduce labour risk, and keep project timelines intact.

Supply chain predictability remains another major challenge the industry faces. Doyle advocated strongly for early engagement with suppliers, noting that proactive collaboration allows developers to secure critical components earlier in the project lifecycle.

By maintaining inventory readiness and reducing lastminute surprises on site, data centre teams can avoid costly schedule disruptions.

Wesco Anixter has shown a strong commitment to enabling efficiency through innovation. From errorproofing lugs to protective cable glands, their



solutions ecosystem is designed to support the pace and precision demanded by modern hyperscale, MTDC, and colocation environments.

Doyle emphasised that in today’s data centre market, “speed is the currency,” and strong partnerships are essential to delivering Already infrastructure on time and at scale.

Wesco Anixter has shown a strong commitment to enabling efficiency through innovation. From error proofing lugs to protective cable glands, their solutions ecosystem is designed to support the pace and precision demanded by modern hyperscale, MTDC, and colocation environments



Redefining the journey: tech-agnostic microgrids for data centres



At AVK, we believe resilience comes from choice. No single technology can power tomorrow's data alone. That's why we design solutions across the full power chain, from turbines and engines to batteries, renewables and future fuels. Our technologyagnostic approach ensures data centres can thrive in a changing energy landscape, as well as providing our clients with the bespoke designs that ensure their data centre meets regulations and requirements.

BY BEN PRITCHARD AVK'S CEO

Introduction: beyond onesizefitsall

DATA CENTRES have rapidly become the nervous system of the digital economy. As artificial intelligence (AI), highperformance computing and cloud services surge, the power profile of the modern facility is changing just as quickly. Loads are larger, power densities are higher, and AI training can drive sharp step changes and long duty cycles. At the same time, grid connections across the UK and Europe are difficult to secure on the timelines the industry needs.

In this context, the old "onesizefitsall" playbook (typically a single technology with a uniform design replicated from site to site) breaks down. A dieselonly strategy may satisfy a narrow standby

brief but struggles with sustainability and permitting. A dual fuel/gas engineonly or turbineonly stance can leave efficiency on the table at low load or hamper agility during AI ramp events. A 'batteriesolveeverything' mindset ignores the realities of longduration supply, fault current, and blackstart needs.

AVKSEG's view is straightforward: resilience comes from choice. Technology agnosticism is not indecision; it's a design philosophy that puts outcomes first and keeps options open – across fuels, OEMs (original equipment manufacturers), and control strategies – so each site gets the right blend for its workload, location and growth path.

Microgrids are the innovative technology that can benefit from our agnosticism.

The development of a microgrid can allow data centre operators to become operational that much faster, providing the solution to today's particular concern of lack of power availability.

This combined with AVK's technological agnosticism can allow for a bespoke data centre design that caters to the sites precise requirements.

Why technological agnosticism matters

Agnostic design avoids the trap of choosing a single "winner" of a solution and then forcing every problem to fit it. Instead, we start with the outcomes –

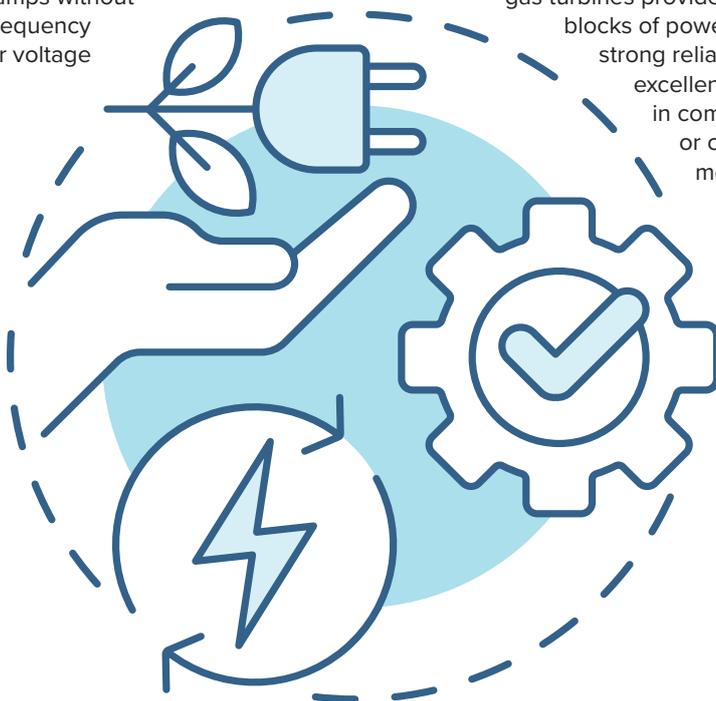
resilience, efficiency, sustainability and cost – and assemble the right toolset for each.

- Resilience. No single technology guarantees uptime under all conditions. Blending generation and storage provides redundancy at multiple layers and avoids single points of failure.
- Efficiency. Different technologies have different sweet spots. Matching assets to the load curve, especially with AI's peaks and troughs, keeps each unit near optimal efficiency.
- Sustainability. The energy transition is uneven. Some regions will prioritise hydrogen blends quickly; others will lead with biogas or aggressive renewables buildout. An agnostic scheme can pivot as local fuel pathways mature.
- Economics. Fuel prices, carbon costs and policy incentives move. A diversified portfolio allows operators to dispatch the lowestcost, lowestcarbon mix in real time.

Load profiles are changing: AI and lowload strategy

There are two practical realities that now shape microgrid design: AI workloads creating spiky demand, and the inevitability of low-load phases.

GPU clusters can drive rapid step changes as AI training jobs start or migrate between nodes. Systems must ride through these ramps without frequency or voltage



excursions. That means fast assets (engines, batteries, gridforming inverters) need to work in collaboration with slower, high-efficiency plants.

Secondly, it is inevitable that there will be lowload phases. Most campuses are built in phases – early halls run at partial utilisation for months, sometimes years. Turbines may operate outside their efficiency sweet spot at low load; engines can cycle inefficiently if not managed; diesel sets wetstack if idled.

A credible lowload strategy is essential: run fewer units harder, use BESS to trim peaks and fill troughs, divert surplus heat into useful cooling, and stage capacity additions to match demand. This strategy is essential to achieve the outcomes outlined above.

The technology toolbox: Setting up for success

There are three main components when setting up power solutions for data centres: the technology, the controls systems, and the mixture of fuels and power. It is the strategic, personalised combination of these components that will provide the most optimal power solution for a data centre. The technology options are vast, and each brings their own design benefits and limitations:

Gas turbines – stable, efficient baseline

Modern aeroderivative and industrial gas turbines provide large blocks of power with strong reliability and excellent efficiency in combined-cycle or cogeneration modes.

Where they shine

- High output density for hyperscale sites with limited footprint.
- Best-in-class efficiency when paired with heatrecovery (steam generation or absorption chilling).
- Lower local emissions per MWh than simple-cycle alternatives at scale.
- Fuel flexibility with pathways to hydrogen blends and, in time, higher hydrogen fractions.

Design notes

- Turbines have turndown limits; heat rate worsens at low load. Pair with engines and BESS so turbines stay in an efficient band while other assets do the “agility work”.
- If water is scarce or permitting is tight, favour aircooled condensers or simple-cycle plus heat-to-chill strategies over traditional steam cycles.

Gas engines – agility, modularity and partload efficiency

Medium-speed reciprocating engines offer rapid start, strong partload performance and modular scalability.

Where they shine

- Fast starts and high ramp rates for AI step loads and spinning reserve.
- High efficiency across a wide load range—ideal for early phases and variable demand.
- Modularity allows “rightsizing” the running set to today's load.

Design notes

- Engines provide valuable fault current and inertia in islanded operation, simplifying protection.
- Plan acoustic treatment early; engine farms in urban locations face tight noise limits.
- Fuel transitions (biogas, synthetic methane, hydrogen blends) are a realistic near-term decarbonisation path.

Battery energy storage systems (BESS) – instant stability and black start

Utility-scale batteries bring capabilities no rotating machine can match.

Where they shine

- Millisecond response to arrest frequency excursions during load steps or generator trips.
- Gridforming operation to provide virtual inertia, voltage control and

Load smoothing to keep turbines and engines in stable operating regions is one of the strengths of BESS. Gridforming operation provides virtual inertia, voltage control and fast fault ride-through, enabling instant stability

fast fault ride-through.

- Black start of auxiliaries and first-up generation after a complete outage.
- Load smoothing to keep turbines and engines in stable operating regions.

Design notes

- Size BESS for power (MW) to handle ramp rates, and for energy (MWh) to ride through short disturbances or support renewable smoothing.
- Integrate with the UPS strategy; avoid duplicated storage by coordinating DC-side and AC-side reserves where appropriate.
- Include lifecycle planning for augmentation, replacement and recycling.

Renewables – lower carbon and operating cost

Onsite solar PV and, where feasible, wind reduce net emissions and hedge against market volatility.

Where they shine

- Near-zero marginal cost once installed.
- Direct emissions reduction without reliance on certificates alone.
- Demand charge management in markets that penalise peaks.

Design notes

- Renewables need firming. Pair with BESS and engines/turbines for smooth integration.
- Consider behind-the-meter PPAs with adjacent renewable assets when onsite space is limited.

Hydrogen, biogas and synthetic fuels – credible transition pathways

Clean molecules are a practical route to lower carbon without rebuilding the plant.

Where they shine

- Hydrogen blends in turbines and engines as supply improves.
- Biogas or synthetic methane using existing gas infrastructure for a drop in reduction in lifecycle emissions.

Design notes

- Design for fuel flexibility from day one: materials compatibility, metering, safety zoning and ventilation for hydrogen.
- Secure long-term offtake arrangements for biogenic fuels; availability is regional.

Diesel/HVO standby – niche but still useful

While not a prime power solution for most new campuses, modern diesel or HVO sets remain relevant as last-line standby or for specific geographies with limited gas infrastructure.

Design notes

- If used, plan for emissions control (SCR/oxidation catalysts) and loadbanking or BESS-assisted loading to prevent wetstacking in low-load tests.
- HVO improves local air quality metrics and lifecycle carbon versus fossil diesel.

Heat recovery, absorption chilling and district energy – turning waste into value

Microgrids unlock thermal synergies that conventional standby cannot.

Where they shine

- Absorption chillers convert waste heat into chilled water, reducing electrical cooling load—valuable for GPU dense halls.
- District heating/industrial heat exports surplus energy to neighbours, creating community benefit and improving utilisation.
- Thermal storage (chilled water/ice) timeshifts cooling, flattening electrical peaks.

Carbon capture – deep decarbonisation option

Where CO₂ transport and storage infrastructure exists or is planned, post-combustion capture on turbines or engines can materially reduce stack emissions.



Design notes

- Allow space and ductwork “stubs” for future capture units; progress policy and offtake discussions early.

Control and protection: the brain of the microgrid

In combination with the technology, controls must be considered during the design stage to ensure optimal performance. Whatever the mix, the architecture of the controls system makes or breaks the scheme.

- Microgrid controller (MGC).** Coordinates dispatch, islanding, synchronisation, powerquality and blackstart sequences. Implements droop control, setpoint optimisation and contingency logic.
- Gridforming inverters.** Enable stable island operation with virtual inertia, fault ride-through and fast voltage control.
- Protection and selectivity.** Islanded systems need intentional fault current sources and carefully graded protection; engines and synchronous condensers can help.
- Harmonics and power factor.** Highdensity IT loads are largely power-electronic. Specify harmonic limits (e.g. THDv/THDi) and reactive support via inverters or STATCOMlike functions.
- Cybersecurity.** Treat the MGC and SCADA as critical infrastructure: network segmentation, secure protocols, rolebased access, monitored remote connections.
- Testing.** Prove black start, islanding, resync and AI load step response during IST — not just nameplate capacity.

How to combine technologies: proven patterns

It is unwise to rely on one form of technology when designing a power solution. The combination of fuel and power types allow for back-up generation, providing reliable power to make sure downtime is non-existent. The proven, successful patterns in technology combinations are as follows:

- Turbine + engines + BESS (hyperscale baseload, AI peaks).** Turbines carry the steady core load efficiently. Engines handle ramps and provide spinning reserve. BESS manages instant transients and provides black start. Add heat



recovery to drive absorption chilling for highdensity halls.

- Engines + BESS + renewables (phased growth, gridconstrained sites).** Engines scale in modular blocks as the campus grows. BESS smooths and provides powerquality. PV/wind offsets daytime demand and lowers OPEX. Ideal where early phases run at lower utilisation.
- Engineonly with BESS (edge or waterconstrained).** Where water for steam cycles is limited or planning favours smaller footprints, an enginecentric scheme offers agility, strong partload efficiency and simpler permitting.
- Turbines + carbon capture + thermal integration (industrial clusters).** In regions with CO₂ infrastructure and heat users, combinedcycle with capture plus district energy delivers very low operational carbon and high overall efficiency.
- Hybrid with HVO standby (strict urban airquality limits).** Gas engines or turbines for prime power, BESS for quality and resilience, and HVOfuelled standby narrowed to rare events—balancing reliability with local emissions constraints.

Designing for geography: legislation and local reality

Agnosticism really earns its keep when policy and permitting differ by country and city. We have diligently expanded our reach throughout Europe, and in doing so it has been important to factor in these differences in legislation to

ensure beneficial outcomes for our clients.

- United Kingdom.** Compliance with G99 (grid code) and the Medium Combustion Plant Directive/Industrial Emissions Directive (as implemented in UK law) drives emissions and monitoring obligations for prime plant. Airquality (AQMA) status will influence abatement, stack height and operating hours. Planning conditions often prioritise noise, visual impact and traffic during construction.
- European Union.** The Industrial Emissions Directive and national transpositions govern thresholds, BAT conclusions and permitting cycles. Fitfor55 measures and ETS costs push design towards higher efficiency and cleaner fuels.
- Local incentives and markets.** Capacity markets, flexibility services, guarantees of origin and green hydrogen pilots vary widely. An agnostic scheme can monetise these where present without depending on them where they are not.
- Infrastructure realities.** Gas availability, CO₂ pipeline proximity, water constraints, and solar resource all shape the optimal mix.

The practical takeaway: define a policyaware reference design for each region, then tailor per site.

The reference ensures repeatability and speed; the tailoring preserves performance and permitting success.

The next decade will reward data centre operators who treat power not as a constraint but as a strategic capability. There is no single technology that can deliver absolute resilience, leading efficiency, and credible decarbonisation across every geography and every workload. But the right combination can. That combination will differ by site, by phase, and over time.

A credible lowload playbook

The early phases are often where inflexible designs stumble. A robust lowload strategy is important, minimising energy consumption and operational costs while enhancing system reliability and sustainability. A low-load strategy typically includes:

- Staging capacity. Commission fewer prime units initially; use modular engines to “rightsize” running plant and add turbines later if baseload stabilises.
- Run fewer, run harder. Avoid idling multiple machines at low efficiency. Keep a small number of units near their sweet spot; let BESS handle short spikes.
- Thermal utilisation. Use absorption chilling or thermal storage so recovered heat remains productive even at modest electrical load.
- Smart testing. Use BESS-assisted

testing to avoid wetstacking on diesel standby and to minimise waste fuel burn.

- Controls tuning. Retune droop, ramp limits and setpoints as each phase adds capacity; today’s parameters will not be right tomorrow.

Procurement without lockin

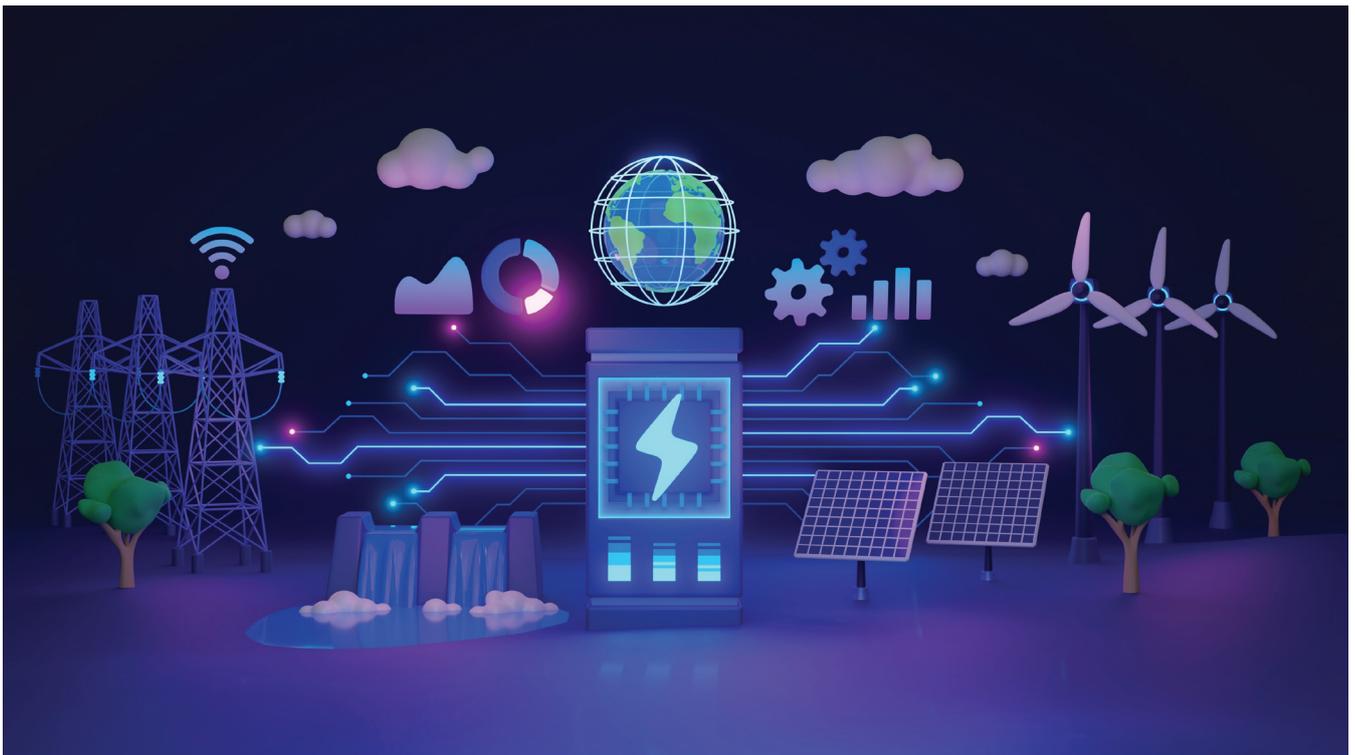
Our agnosticism in technology also extends to our supply chain relationships and contracts. This agnostic approach is seen by prioritising multi-OEM optionality, open interfaces, performance-based outcomes, and lifecycle clarity.

This allows us to de-risk manufacturing slots, integrate suppliers seamlessly, tailor contracts for the best outcomes and plan for the future. Above all else, our agnosticism is one of the reasons for our high rate of client satisfaction.

Conclusion: flexibility is the advantage

The next decade will reward data centre operators who treat power not as a constraint but as a strategic capability. There is no single technology that can deliver absolute resilience, leading efficiency, and credible decarbonisation across every geography and every workload. But the right combination can. That combination will differ by site, by phase, and over time.

Technology agnosticism is how you keep that combination optimal. It’s how you accommodate AI’s volatile demand without overbuilding, how you navigate permitting from Dublin to Düsseldorf, and how you migrate to cleaner fuels without stranding assets. It is, in short, how you build microgrids that are as dynamic as the digital world they power.



Future-proofing data centres through on-site power and distributed carbon capture



There is finally a pathway for data centres to address both power needs and decarbonization, while maintaining full control over power supply - without compromising reliability, cost, or sustainability.

BY ANNA PAVLOVA, SENIOR VICE PRESIDENT FOR STRATEGY, MARKET DEVELOPMENT, AND SUSTAINABILITY, CARBONQUEST

THE PAST FEW YEARS have marked a time of unprecedented growth for data centers. Across primary U.S. markets, [the number of data centers surged](#) in 2024, including rapid development beyond traditional markets and into rural areas in Appalachia, Western states, and the South. While this growth has been necessary to facilitate the AI boom, behind this momentum lies a pressing concern — the growing power demands of these facilities competing for electricity, combined with the often contradictory sustainability demands of customers.

The Rocky Mountain Institute 2024 report, "[Powering the Data Center Boom with Low-Carbon Solutions](#)", highlights that most data center hubs are either running out of grid capacity or facing challenges supporting the electricity consumption growth led by data centers.

Data centers must rapidly find solutions beyond traditional efficiency and cooling improvements. But utilities and the transmission infrastructure are not set up to bring online clients that seek gigawatts of power, especially when it comes to clean power with a 99.9% capacity factor. Hence, we see the re-emergence of nuclear power, agreements outside of traditional utility structure, and renewed interest in onsite power generation. In fact, [a recent survey](#) by Bloom Energy shows that 30% of data centers are planning to use onsite power as a primary source of electricity.

Combined Heat and Power (CHP) systems and fuel cells provide reliable onsite power that can operate independently of the grid. Additionally,

because CHP systems convert chemical energy into electrical energy with fewer losses, these technologies offer attractive onsite energy efficiency. However, whether such a system consists of fuel cells, RECI engines, or turbines, the primary power generation source is usually natural gas. While CHP solves the problem for reliability and onsite power generation, it still faces an obstacle where the data center developer requires low or zero emissions.

To address this challenge, there's no better option than integrating CHP and fuel cell systems with distributed, modular, and stackable carbon capture technologies. By [combining distributed carbon capture \(DCCS™\) with alternative power sources](#), data center operators can achieve significant reductions in greenhouse gas (GHG) emissions, thereby aligning with global climate objectives and the commitments of key customers, while addressing electricity consumption needs onsite.

Distributed carbon capture technology is designed to operate concurrently with CHP or fuel cells, trap the CO₂ before it is emitted into the atmosphere, and turn the captured CO₂ into a useful industrial product, mineralize it into rock, or sequester it underground.

Distributed carbon capture solutions provide multiple advantages over the traditional stationary amine-based carbon capture systems. This type of

technology does not require chemical solvents that create toxic water challenges, or 150-foot-tall towers that require height permitting. The systems use solid sorbents to trap the CO₂ in a series of stackable 6-foot vessels.

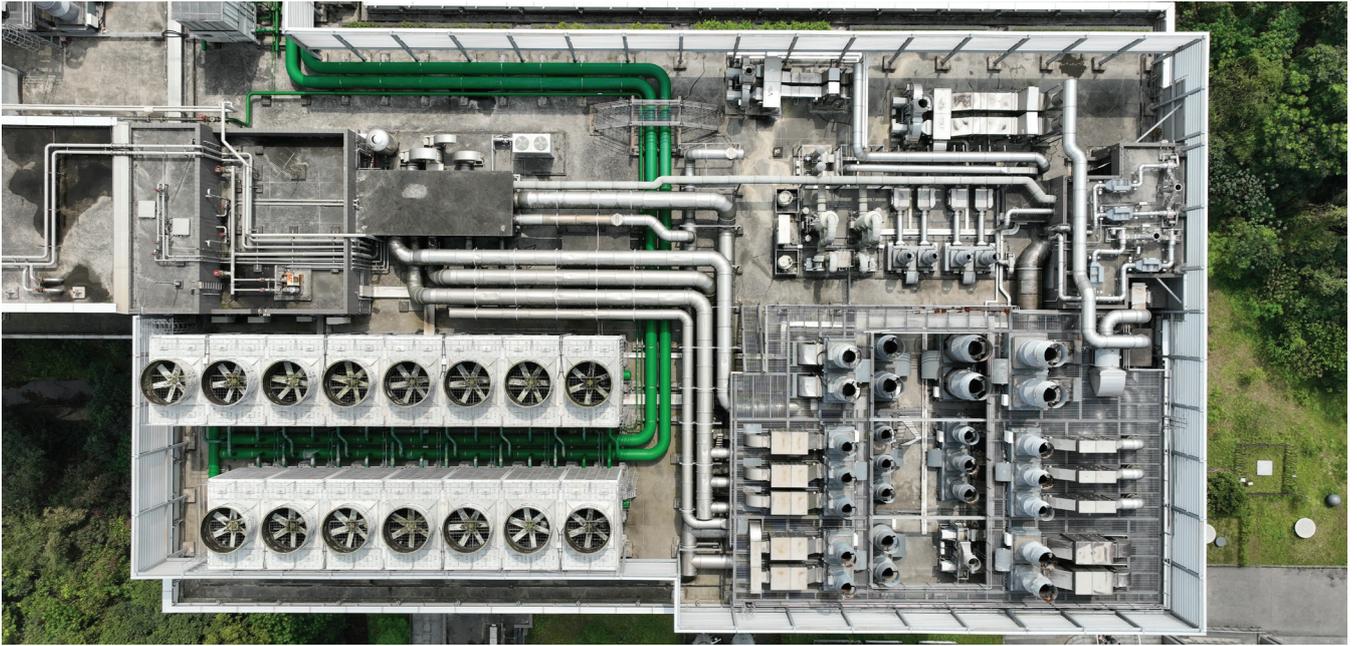
Additionally, because these systems are containerized, fully modular, and stackable, they can easily be expanded as the needs of data center facilities grow by adding additional trains. Finally, the systems are already pre-designed for use with different types of fuel cells and engines, with the entire package offering a cost-competitive solution in virtually any geographic region.

The synergistic integration of onsite power and carbon capture enables data center operators to future-proof their facilities against regulatory uncertainties and customer pressures for sustainability, while ensuring a predictable supply of high-capacity factor, automated onsite power supply.

This solution can be combined with other options like small nuclear reactors, utility power, and renewable energy power purchase agreements.

There is finally a pathway for data centers to address both power needs and decarbonization, while maintaining full control over power supply - without compromising reliability, cost, or sustainability.





Beyond cooling: how data centres can become heat hubs for Net Zero infrastructure



Every year, data centres in London release enough waste heat to warm up to half a million homes, according to AECOM and the Greater London Authority. Most of it is vented straight into the atmosphere. But buried beneath headlines about megawatt-hungry facilities is an overlooked reality: data centres don't just consume energy - they also produce it.

BY CHRIS DAVIDSON, CTO AT GENIUS ENERGY LAB

EVERY WATT OF POWER that runs a server eventually becomes heat, and with the right approach, that heat doesn't have to be wasted. It can be captured, upgraded, and repurposed to warm nearby homes, schools, hospitals and businesses.

Heat reuse projects are already underway in London and Tallaght, backed by public funding and growing regulatory support. For developers, that opens up a new strategic consideration: where and when could heat reuse add value to the community, the planning case, or the bottom line?

The missed opportunity in energy planning

Despite the scale of energy flowing through data centres, heat reuse is rarely considered at the design stage. Planning conversations typically focus

on electrical load and cooling efficiency, while the thermal energy being ejected from the site is overlooked.

That's a missed opportunity - data centres produce a steady stream of low-grade heat 24/7 - and with rising gas prices and growing pressure to electrify heat, that waste is becoming harder to justify. Treating it as a usable resource, rather than an unfortunate by-product, could reduce costs, improve resilience and create tangible benefits, both for operators and the communities they operate in.

The technology is already available. Systems like heat exchangers, ground source heat pump (GSHP) boreholes and district heating networks are proven and in use - including in data centre settings. These are the tools that can capture, store, upgrade and

distribute heat locally. GSHPs are especially well-suited here: they're highly efficient, can be scaled for large developments, and enable inter-seasonal heat storage in the ground.

What's missing is integration - developers, utilities and local authorities often work in silos, with limited coordination around thermal infrastructure. This means that viable reuse opportunities often go unnoticed.

It's also a question of mindset. Data centres are still often treated as standalone assets - consuming energy, creating jobs, but disconnected from the communities around them. That model isn't sustainable, in any sense of the word. As demand for low-carbon heat grows, integrated design - where digital and thermal infrastructure are planned together - is a better solution

for both operators and the communities they operate in.

Inside the heat reuse model

So how does heat reuse from data centres actually work?

The principle is straightforward, even if delivery requires careful coordination. Low-grade heat generated during data centre cooling is captured before it's vented. That thermal energy is then transferred via heat exchangers and fed into a ground array - typically a set of vertical boreholes drilled into the earth beneath or near the site.

The ground acts as a thermal battery, storing and balancing heat throughout the year. Ground source heat pumps (GSHPs) then extract that stored heat, upgrade it to a usable temperature and feed it into a district heating network - supplying homes, schools, hospitals and commercial buildings.

These systems are already operating across Europe, including a live project in Tallaght, Dublin, where data centre heat is captured, upgraded and piped into a local district network. In the UK, a large-scale scheme is now underway in West London that will recover data centre heat for use in a new district heating system serving homes and commercial buildings. Integration is proven and practical, especially where local demand is high and the network is within reach. The key is early-stage planning. Retrofitting is possible, but co-designing thermal and digital infrastructure from the outset is cheaper, more efficient and far more scalable.

This model also provides flexibility, as stored heat can be drawn on when it's needed, rather than just when the data centre is producing excess. That means the system can smooth demand, reduce grid stress and operate more like a real energy asset than a passive donor.

Why now?

Until recently, the idea of reusing waste heat from data centres sat on the fringes of energy strategy - a niche opportunity, often dismissed as too

complex or too site-specific to scale - but that's starting to change. Several factors are converging to bring heat reuse into the mainstream.

First, governments are beginning to actively support it. In the UK, funding from the Green Heat Network Fund is enabling large-scale projects, including the West London scheme. In Ireland, the Tallaght system has become a case study for how data centre heat can reliably serve public buildings. The UK government is also preparing legislation on Heat Network Zoning, which would mandate certain buildings - including new developments - to connect to local heat networks where available.

Second, planning requirements are evolving. In some regions, local authorities are starting to ask for clearer answers on how new developments will contribute to net zero goals, including what happens to waste heat. For data centre operators, being able to demonstrate a heat reuse strategy can support planning approvals and build goodwill with local stakeholders.

Finally, energy economics are shifting. The drive to decarbonise heating - one of the most difficult parts of the net zero equation - is creating new urgency around viable alternatives to gas. Data centre heat may not be a universal solution, but where the conditions are right, it offers a scalable, low-carbon supply.

The business case

While often framed as a sustainability initiative, heat reuse has real commercial implications for operators.

A credible heat reuse plan can support smoother planning processes, especially in urban or sensitive locations where data centre developments often face public or political resistance.

It also contributes directly to ESG goals, offering a tangible local impact that's increasingly important to investors, customers and regulators.



Beyond reputational benefits, a well-planned heat reuse strategy can strengthen the overall case for development. It may help reduce infrastructure duplication, demonstrate alignment with local energy priorities, or open up future collaboration with public bodies and utilities. For example, ground arrays typically last around 100 years, making them ideal for futureproofing local heat infrastructure. Even where implementation is complex, the ability to show forward thinking on heat can be a meaningful differentiator.

Of course, not every site will be suitable for heat reuse, and not every opportunity will stack up commercially. With the right local conditions, partners and design strategy, it becomes a credible addition to the planning and sustainability case for new developments.

Putting heat on the agenda

As energy systems become more local, decarbonised and interconnected, the boundaries between digital and thermal infrastructure are blurring - drawing data centres into wider conversations about how we plan and manage heat.

That shift brings complexity, but also opportunity. Operators who think ahead and engage early with local heat strategy will be better placed to navigate planning challenges, demonstrate public value and align with long-term energy goals. In a sector facing growing scrutiny, the ability to deliver low-carbon heat alongside digital services could be what sets the next generation of data centres apart.

With the right local conditions, partners and design strategy, heat reuse becomes a credible addition to the planning and sustainability case for new developments

Unhackable by design: securing AI data centres at the physical layer



AI data centres are fast becoming the backbone of the digital economy. They process the most sensitive data, power near-constant workloads and underpin critical services. That makes them indispensable and uniquely exposed.

BY MICHAEL VALLAS, GLOBAL TECHNICAL PRINCIPAL, GOLDILOCK SECURE

BY THE END OF 2025, over a third of global data centre capacity is expected to be dedicated to AI workloads. According to McKinsey, demand for capacity overall is projected to grow by more than 20% each year through 2030. As these environments scale, the attack surface grows with them.

Colocation and hybrid models magnify the risks further, creating more entry points and more opportunity for attackers to move laterally. It's happening now, with cybercriminals and nation-state actors already zeroing in on these weaknesses. The UK government's decision to designate data centres as part of the nation's

critical infrastructure underscores just how high the stakes have become.

And yet, protection models remain skewed toward software-first strategies: cyber defences represent vast and potentially fragile lines of code trying to keep pace with automated, adaptive attacks. The old approach feels a bit like building a firewall in a burning forest: it's missing the bigger picture.

The limits of software-first security

Almost every cybersecurity battle today is still fought in code. We patch, configure and layer on new tools, and attackers respond by finding new

cracks. The result is an endless cycle of software versus software, with defenders overstretched.

Firewalls and endpoint security remain vital parts of a layered strategy, of course. But like all software-based tools, they have their limitations. Even a brief compromise in critical environments can disrupt essential services or expose sensitive AI-driven data. When detection means sifting for faint signals across trillions of daily events, delays are inevitable and dangerous.

The case for physical resilience

True resilience means gaining control over the physical pathways that carry



data in and out, and the networks that connect critical systems. This is where hardware-enforced isolation comes in.

Physical isolation allows operators to instantly disconnect compute, storage and network segments with secure, out-of-band commands that sit outside the attack surface.

The concept is simple but powerful (or powerful because it's simple): if malware can't reach the system, it can't compromise it. And unlike software-only controls, physical isolation can't be tampered with remotely. There's no IP address, no hypervisor dependence, no accessible or exploitable code: just a clean physical break.

Critically, this doesn't mean downtime. Systems can continue running safely in an offline state, maintaining core operations while remaining unreachable to attackers. Organisations can decide when to be connected and when to disconnect, moving from an "always-on" mindset to a risk-aware, resilient model.

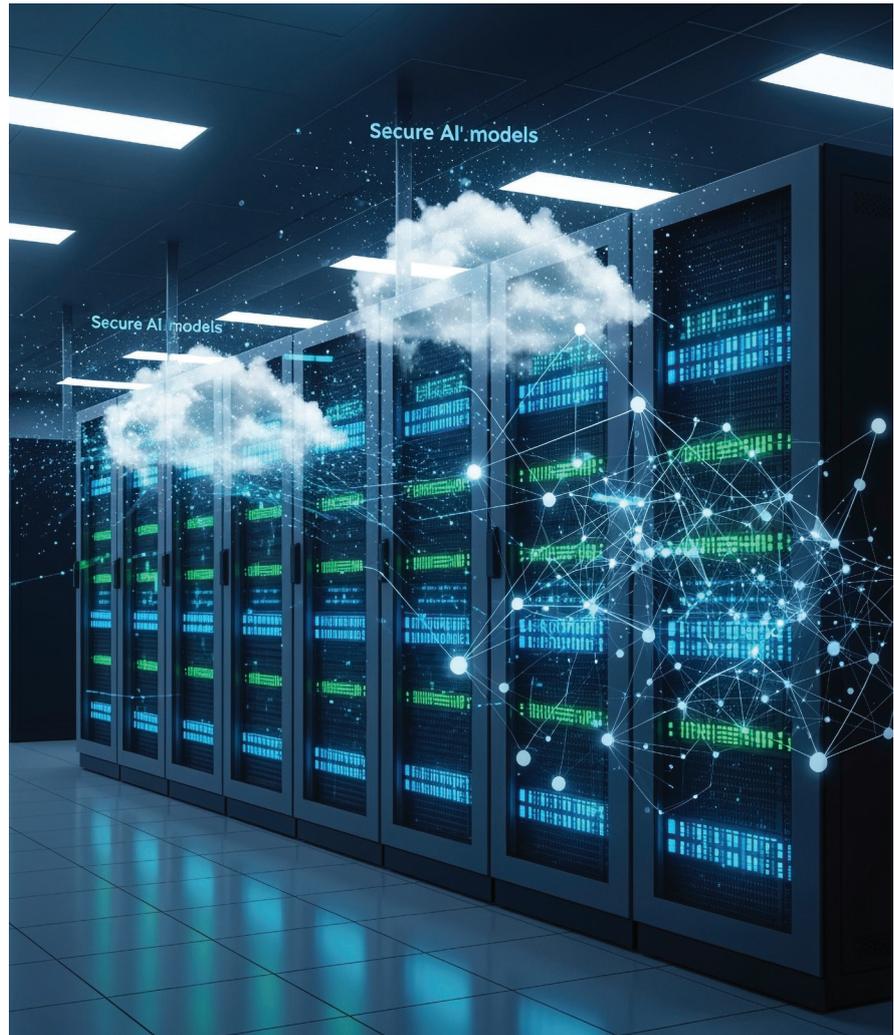
Where physical isolation matters most

The value of isolation is clearest in high-stakes environments where speed and certainty matter. In colocation facilities, it prevents cross-tenant spread by cutting off a compromised segment before the threat can move laterally.

For enterprise IT, critical administrative systems can be isolated during high-risk operations or when threats are detected, containing potential damage while keeping core business functions running.

At disaster recovery sites, systems can remain physically offline until needed, ensuring clean, uncompromised backups are always available to restore services.

In cloud and backup environments, selective disconnection ensures ransomware cannot encrypt critical archives. And across AI-heavy workloads, hardware isolation blocks



data exfiltration and model tampering, while enforcing strict security boundaries around sensitive processes.

The bottom line

As AI becomes embedded in everything from healthcare diagnostics to financial systems to national security, the infrastructure behind it must be absolutely trusted. And here's the thing: complete trust doesn't come from adding more layers of software. It comes from designing resilience into the system itself.

The NCSC has already urged organisations to build in the capability to fully disconnect critical systems from networks: a clear signal that such measures could soon become

regulatory requirements. Taking that step not only strengthens defences today but also positions organisations ahead of potential regulation tomorrow.

All of this comes down to protecting the backbone that modern society increasingly depends on. When AI systems control everything from power grids to medical devices, a single breach can threaten public safety, disrupt economies and undermine national security.

In the end, it comes down to this: do we only keep chasing patches in a game we can't win, or do we also build a defence designed to hold no matter how the threat evolves?

The concept is simple but powerful (or powerful because it's simple): if malware can't reach the system, it can't compromise it. And unlike software-only controls, physical isolation can't be tampered with remotely



The firewall: mitigating data centre fires with stone wool sandwich panels



In a time when digital continuity directly impacts public safety, economic stability, and social interaction, we must rigorously audit the backbone of our digital infrastructure. Whether from conventional hazards or next-generation upgrades like rooftop solar, fires can undo sustainability progress and disrupt millions of lives.

**BY DANIELA PASQUERO PHD, PUBLIC AFFAIRS AND INNOVATION MANAGER
AT ROCKWOOL CORE SOLUTIONS**

WE ARE IN A digital era defined by relentless data generation. Whether processing financial transactions, streaming videos, managing smart buildings, or training artificial intelligence models, data centres sit at the epicentre of our modern lives.

Global spending on data centres is projected to surpass \$1 trillion in the coming years, driven by exponential growth in cloud computing, AI, and digital transformation¹.

However, as data centres flourish both in scale and mission-critical importance, a vital question must be asked: Are we building them safely and sustainably enough to withstand the risks of today and tomorrow?

The fire risk facing the digital economy

Recent data centre fires have highlighted the potential fragility of even the most advanced facilities. In March 2021, a major fire at the OVHCloud facility in Strasbourg destroyed one building and rendered another inoperable. Millions of websites, including government, banking and news platforms, were disrupted. Similarly, in March 2023, a fire triggered by a water leak and subsequent UPS failure led to a multi-cluster shutdown at Google Cloud's facility in Paris, severely affecting over 90 cloud services.

These aren't isolated incidents. According to the Uptime Institute, over

60% of data centre outages cost more than \$100,000, with many exceeding \$1 million². The average cost of data centre downtime now stands at roughly \$9,000 per minute³. Beyond the financial toll, outages directly impact public trust and operational continuity for businesses and institutions alike.

Mitigating this risk begins with smarter design. A critical line of defence lies in passive fire protection built directly into the structure, especially the building envelope. Enter non-combustible stone wool sandwich panel insulation.

Prioritising fire resistance

Stone wool, made from molten volcanic rock, is inherently non-combustible, withstanding temperatures above

1,000°C. When used as the insulating core in sandwich panels, it creates a robust fire-resistant barrier within the building façade, and partition walls.

Classified A2-s1, d0 under EN 13501 standards, stone wool sandwich panels do not propagate flames or emit toxic smoke. This crucial feature buys valuable time for evacuation, emergency response and system failover, minimising both damage and danger. These panels maintain structural integrity under fire, slowing the spread and compartmentalising risk.

This passive protection is essential in a world where emerging technologies introduce new fire threats: from lithium-ion batteries to high-intensity computing loads to, most notably, rooftop solar installations.

Sustainable intent; managed risk

As part of sustainability commitments, data centre operators are increasingly deploying rooftop solar photovoltaic (PV) arrays. While this transition to renewable energy is commendable, it introduces unintended risk.

Solar panel fires are on the rise. A 2020 study by BRE National Solar Centre in the UK found that fire incidents involving PV systems, often due to faulty connectors, poor installations,

or panel malfunctions, are more common than previously thought⁴. Fires originating from faulty rooftop solar installations can rapidly spread across combustible roofing elements, threatening the entire facility and all the data it protects.

Making a data centre 'green' should not introduce hazardous vulnerabilities. We must ensure sustainability drives are themselves sustainable in every dimension, including fire resilience.

Energy performance and thermal efficiency

Beyond fire safety, stone wool-enhanced sandwich panels provide substantial energy efficiency benefits, which are especially critical given that data centres currently consume 2% of global electricity; a figure expected to rise to 9% by 2030⁵.

High-density stone wool insulation reduces thermal transfer, enhancing the stability of internal environments while lowering the demand for active HVAC solutions. As thermal displacement is limited, cooling loads are reduced, which is vital in high-density racks consuming upwards of 60kW each due to AI workloads⁶. These panels also support green roofing systems, aiding cooling and stormwater management while mitigating the urban heat island effect.

High-density stone wool insulation reduces thermal transfer, enhancing the stability of internal environments while lowering the demand for active HVAC solutions

In the pursuit of net-zero or carbon-neutral operations backed by solar, wind, and battery storage, sandwich panels with stone wool insulation align with green design goals while maintaining operational integrity.

Modularity and the construction advantage

Modern data centres must not only be sustainable but scalable. Here, sandwich panel systems support industrialised, off-site construction approaches that favour speed, precision, and adaptability.

Dry construction systems using prefabricated sandwich panels allow architects and engineers to rapidly deploy high-performance building envelopes. Modularity simplifies future expansion as hardware evolves or



leasing needs shift, supporting the 'growth on demand' model more sustainably.

Moreover, stone wool's longevity and recyclability add to the lifecycle value. At end-of-life, steel panel facings and stone wool cores can be recovered and reused, contributing to circular design principles.

Reactivity to resilience

Clearly, the best fire is the one that never spreads. In addition to shaping a building's environmental performance, materials selection defines a data centre's passive resilience. When considering the extreme financial penalties of outages, asset loss, destroyed equipment, and reputation damage, the return on investments in non-combustible materials becomes evident.

Data centre operators also face evolving insurance landscapes. Premiums covering data centre facilities — some of which cost over \$1 billion to construct — are rising steadily. Insurers are becoming wary of poorly mitigated fire and downtime risks⁷. Integrating flame-resistant solutions like stone wool can reduce exposure, making facilities more insurable and financially sound.

Safeguarding the vision

In a time when digital continuity directly impacts public safety, economic stability, and social interaction, we must rigorously audit the backbone of our digital infrastructure. Whether from conventional hazards or next-

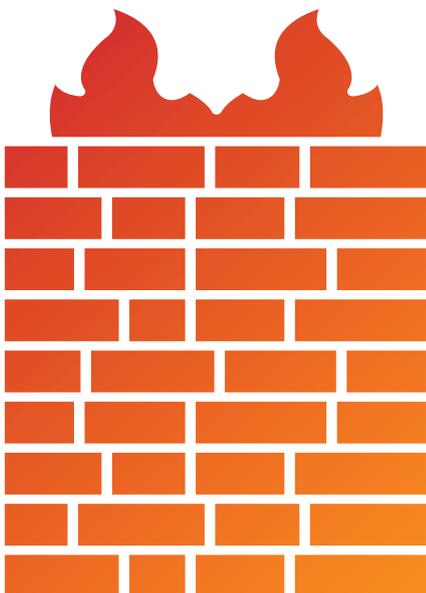


generation upgrades like rooftop solar, fires can undo sustainability progress and disrupt millions of lives.

As architects, engineers, and developers, we have a duty to shift from reactive risk response to proactive resilience. That means designing structures for both durability and flexibility, which are prepared to embed clean energy systems without fire vulnerability.

With stone wool core for sandwich panels, ROCKWOOL Core Solutions supplies an indispensable tool for enabling this change.

They provide thermal performance, acoustic comfort, structural flexibility, and most vitally, unrivalled passive fire protection — without compromising sustainability. For data centres tasked with delivering 24/7 uptime in a carbon-constrained future, this is not a compromise. It's a new standard.



SOURCES

- [1] CBRE Data Centre Trends Report, 2023
- [2] Uptime Institute Annual Global Data Center Survey, 2023
- [3] Ponemon Institute: Cost of Data Center Outages
- [4] BRE National Solar Centre: Fire and Solar PV Systems Report, 2020
- [5] International Energy Agency (IEA), Data Centre Energy Report, 2023
- [6] Schneider Electric: AI and High-Density Rack Trends, 2024
- [7] Marsh: Construction Insurance Insights — Data Centres Report, 2024



New blog from AFL

Hyperscale Market Shifts: AI, Neoclouds, and the New Limits of Data Centers

AI is driving a new era of hyperscale data centers built for continuous, large-scale workloads. Rack-scale GPUs and high-speed fabrics (400G, 800G, and beyond) are accelerating changes in cluster size, topology, and fiber density.

Read AFL's blog: [Hyperscale Market Shifts: AI, Neoclouds, and the New Limits of Data Centers](#). See how AI is influencing hyperscale infrastructure, and what these shifts mean for fiber, network design, and physical planning.

Read the blog now:

[Market Shifts: AI, Neoclouds, and the New Limits of Data Centers](#)

<https://www.aflhyperscale.com/articles/hyperscale-market-shifts-ai-neoclouds-and-the-new-limits-of-data-centers/>

The five Ss: a guide to selecting effective data centre fire suppression



The modern data centre may not always be a single, large space and may instead be a series of data halls. The heat generated by the servers in these halls poses a fire risk, so protection methods should be selected accordingly. Good fire system design can help data centREs protect people, equipment, uptime and business continuity in the event of a fire.

BY ALAN ELDER, FELLOW — INDUSTRY RELATIONS, JOHNSON CONTROLS

THERE ARE SEVERAL firefighting options suited to data centers:

- Sprinkler systems
- Water mist systems
- Halocarbon gas extinguishing systems
- Inert gas extinguishing systems

However, there is no one-size-fits-all solution. Five critical factors (safety, space, speed, spend and sustainability), referred to as the “Five Ss”, can help prioritize protection objectives and guide system design. Considering these Five Ss can help determine the best firefighting solution for each data center.

1. Safety

In most data centers, protecting occupants and maintaining system uptime are typically the highest priorities. Therefore, a key protection objective is to extinguish a fire and prevent its spread using an agent that’s safe for people and equipment.

Gas systems using halocarbons and inert gases are designed primarily to extinguish a fire. This discharge of a gas extinguishing system, especially in a non-fire

condition, will not adversely affect people or equipment. Water-based systems, such as sprinkler systems and water mist systems, are engineered to control or suppress a fire¹. The discharge of a water-based system will not harm people (other than getting them wet) but may damage equipment.

In addition to these systems, aerosols have also been used for firefighting in data centers. If aerosols discharge, the particles may be inhaled by occupants, so these systems should not be used

in occupied areas. Aerosols also leave a fine

residue that can damage equipment and require extensive cleanup. European Standard EN 50600-2-5 advises that aerosols should not be used in occupied spaces or in spaces containing electronic equipment.

Overall, a gaseous system backed by a water-based system provides an optimal level of protection without



- Acoustic nozzles reduce the sound pressure of a typical gas suppression system discharge to help protect sensitive data center equipment.



➤ Good fire system design can help data centers protect people, equipment, uptime and business continuity in the event of a fire.

undue risk of damage to equipment or the safety of people.

2. Space

Modern data centers typically have a room allocated specifically for fire protection system equipment. The space required for this equipment depends on how much room is available and the type of fire protection equipment required in the region where the data center is located. Some regions — including North America — mandate a sprinkler system for building protection. In that case, if the space is already allocated, other protection objectives may call for the addition of a gas system.

Gaseous systems usually require less space than water sprinkler or water mist systems. Gaseous systems use gases stored under pressure within containers, providing a comparatively large firefighting capability in a relatively small footprint. Sprinklers and water mist systems typically require pumps and a water tank, which may have a larger footprint than gas containers.

3. Speed

To prevent a fire from spreading, it's critical to contain and extinguish it in

its earliest stages. Gas systems make it possible to contain and extinguish a fire within minutes of first detecting smoke. Gas systems are connected to an early warning fire detection system designed to sense smoke in the very early stages. After more than one detector senses smoke, the sequence to deploy the system begins. After a short time delay, usually not more than 30 seconds, the gaseous system begins to discharge. Halocarbon systems discharge entirely within 10 seconds for rapid extinguishment, while inert gas systems for data centers take about two minutes to discharge fully. During the discharge period, the agent impacts the fire until it's extinguished.

With water-based systems, the discharge is dependent on heat. Heat can take time to accumulate during a fire event, especially if there is high air movement in the space, like that from mission-critical HVAC equipment. Depending on the device's location, considerable time may pass before there is sufficient heat to activate the sprinkler head or water mist head. However, pre-action sprinklers connected to a detection system could detect smoke and prime the sprinkler system.

It should be considered that air movements within the protected space can also affect smoke's ability to reach the smoke detector, so it is important to ensure the fire detection system's design accounts for the airflows present.

Gas systems have a faster response compared to water-based systems and can detect fires in incipient stages, making a gas system backed by a water-based system ideal.

4. Spend

In a perfect world, every data center would have an early warning detection system and a gaseous system backed by a water-based system. However, budgets need to be respected, and system design governed by budget starts with what is mandated by code. In North America, for instance, building code requires a sprinkler system. While the size and complexity of the data center will dictate the upfront cost of a firefighting system, there are additional considerations.

Hazard geometry and the cost of cleanup and downtime after system discharge can also impact the total cost of ownership.



➤ For most data centers, protecting occupants and system uptime are typically the highest priorities.

Understanding hazard geometry, including factors such as air movement, can help predict system efficacy during a fire event. For instance, air handling equipment is critical to data center performance and uptime and cannot be shut down. If a gas system discharges while air handling equipment is running, it will not hinder system performance and may even help by drawing agent into affected equipment. In comparison, the high airflows of air-handling equipment can detrimentally affect the performance of a water-based system by moving water away from the fire. The design of the water-based system should recognize the airflow rates and determine whether it can be deployed within the system's approved limits.

The costs of cleanup, downtime and equipment replacement must also be considered. It is very likely that the cost of cleanup and downtime will be significantly higher with a water-based system than with a gas system backed by a water-based system. Gas systems use clean agents that produce

no residue. In the event of discharge, a gas system requires post-discharge ventilation, resulting in minimal cleanup. Downtime is limited to equipment damaged by fire. In the event of a water-based system discharge, affected spaces require cleanup, while equipment damaged by water and/or fire typically requires replacement. There is some downtime until both are complete.

5. Sustainability

Sustainable operation is a high priority for data centers, and environmental impact extends to fire protection. For a water-based system, consumption may be a consideration in water-insecure areas. For these regions, a gas system may be preferred.

There are no environmental concerns for inert gas agents, such as those offered in the Johnson Controls IGS-300 system. And while there may be environmental concerns about halocarbon agents, FK-5-1-12, which is offered in the Johnson Controls

SAPPHIRE system, is a fluoroketone with a negligible global warming potential and is exempt from use restrictions arising from concerns about climate change.

Designing ideal fire protection for each data center

Every data center is unique, and its fire protection solutions should align with its needs and core protection objectives. To streamline the selection process, partner with a provider that offers flexible, fully compliant solutions and proven performance.

They can help you leverage the Five Ss to prioritize objectives and design the ideal fire protection solution.

REFERENCE

- [1] It is possible a water-based system could extinguish the fire, but it is not the design premise.

the lifeblood of every modern data centre



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

scan to explore



Strong foundations: How to get your physical data centre security right from the start



Data centres are, rightly, renowned for their physical as well as digital security. As the number of data centres is set to grow rapidly, by almost a fifth in the UK alone, how can data centre leaders ensure their security infrastructure is consistent across their portfolio, is easy to maintain, and intuitive for operators to monitor?

BY JAMES SMITH, SALES AND MARKETING DIRECTOR - RELIANCE HIGH-TECH

MOREOVER, as the industry scales, infrastructure alone isn't enough. The real differentiator is the expertise behind it, the integrators who design, install, and maintain systems to consistently high standards. Without that human foundation, even the best infrastructure risks leaving cracks in your security posture.

Treat physical security and people processes with the same level of rigour as the data held inside. From careful planning at the start of a build to seamless day-to-day operation, the smallest details matter when creating a best-in-class security system across multiple locations.

Putting in place some physical security

best practices that become second nature at every site development will ensure your people, processes, and infrastructure aren't the weak link that causes unexpected downtime, or worse.

Build for resilience from day one

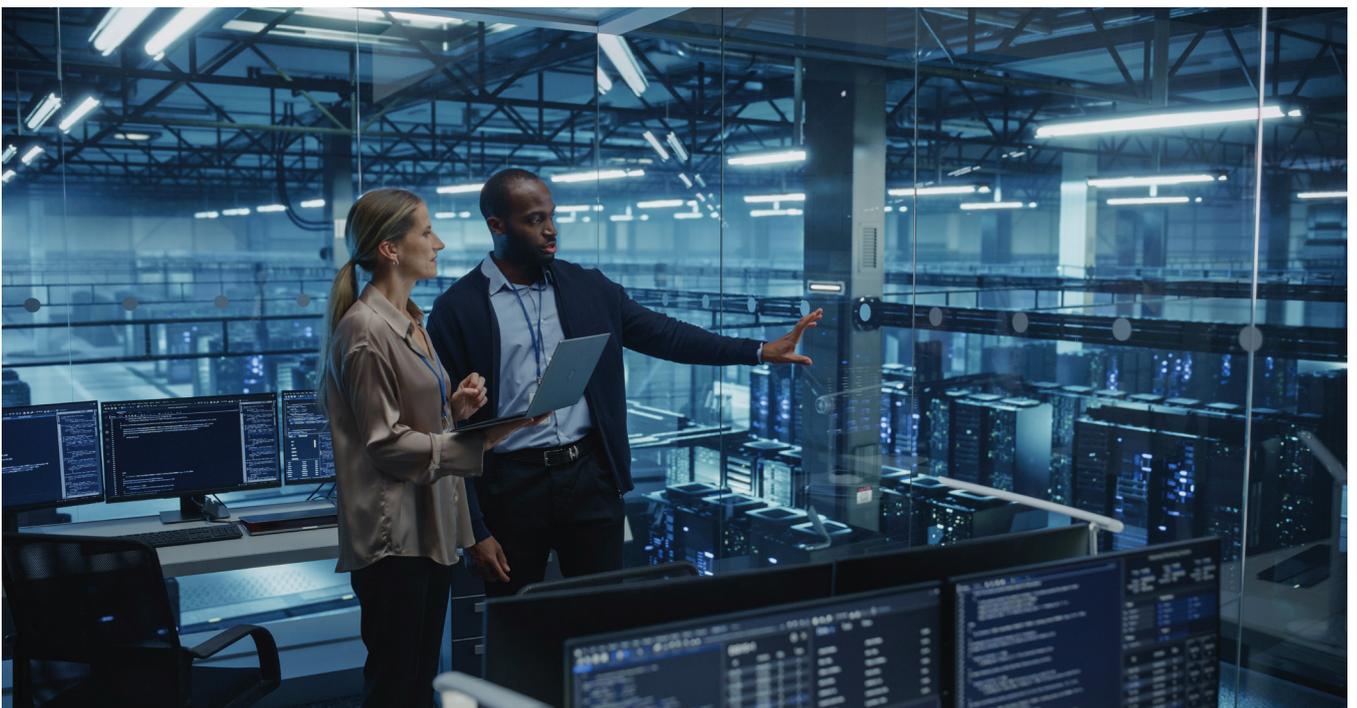
The most successful data centre security projects start with the basics done well. Clear direction from the start, for the entire project team, can make a huge difference to the ultimate success of your security installation.

Appointing a skilled integrator that assigns a dedicated project manager with an understanding both of security and the pressures of an in-build or

live data centre, ensures quality from design through to operation. Backed by the right accreditations, experience and communication skills, the team gets it right first time. Working to repeatable high standards and supported by project engineers who coordinate seamlessly with other trades, for example, delivers the 'hidden benefits' of good people - smoother delivery and consistently high-quality outcomes.

Prioritise people and processes

Even the most advanced security technology is only as strong as the people behind it. If you're lacking the right integration expertise to design, install, and maintain security systems, vulnerabilities could be introduced. Trust, vetting, and accountability are



as vital to your physical security as the cameras, VMS, access control, AI/ analytics, perimeter protection and more.

This rings especially true now that data centres are formally recognised as part of the UK's Critical National Infrastructure (CNI). Leaders need project teams who understand not only the nuances of security technology, but also the stringent standards expected across other CNI environments, where failure isn't an option. Partnering with an integrator who can offer project management and sector expertise will help you keep your security installation running smoothly and on time. The best integrators will have project managers who have walked the line between engineering detail and big-picture delivery so nothing gets 'designed in' by accident. That expertise extends beyond installation. The best integrators invest in their teams and provide ongoing training and maintenance regimes, from patching vulnerabilities to updating systems, so your infrastructure never drifts out of date or develops weak points.

Think beyond a single site

Expansion is almost inevitable, due to demand for data centre capacity, but as your portfolio grows, you want to make sure the complexity of managing and monitoring multiple sites doesn't increase. To avoid this, ask the right questions early on in your security implementation:

- **Integration:** Will my security system work seamlessly with other site technologies, as well as across multiple facilities?

If your security technologies connect seamlessly with other on-site systems, and across your wider portfolio, you create a consolidated view that's easier to monitor and uncover data insights.

- **Usability:** Do operators have a consolidated "single source of truth," or are they juggling multiple systems?

If systems aren't intuitive, operators need more training, mistakes are more likely, and efficiency suffers.

- **Maintainability:** Who is installing the solution, and will they remain accountable for updates, maintenance, and cross-site consistency?

Choosing solutions that are straightforward to update and manage across multiple facilities reduces hassle, keeps total cost of ownership under control, and avoids the inefficiencies of juggling different platforms.

The same high standards need to be replicated across every facility in your portfolio. That only happens when you work with experts who can design, install, and maintain systems reliably every time.

Get outside expertise where needed

You don't have to design or implement security technology alone. Choosing the right integration partner to recommend solutions, install them, manage your project end-to-end, and maintain your system afterwards takes a significant weight off your shoulders, giving you the space to focus on the bigger picture for your data centre.

Leading integrators feel like an extension of your team, safeguarding your interests and ensuring that standards never slip. Their ability to manage multiple vendors while keeping you informed is what prevents gaps and ensures your infrastructure works smoothly day to day. Vendor-agnostic integrators can recommend the best solutions for your needs without the risk of lock-in to a single ecosystem.

When choosing a partner, look for the following trustmarks: accredited in-house engineers, transparent quality controls, and recognised certifications such as Cyber Essentials Plus, ISO 27001, and ISO 22301.

These are green flags that the fundamentals will be delivered properly, without hidden vulnerabilities.



No weaknesses, no compromise

The goal for your data centre security is simple: to protect your assets effectively and efficiently, scaling to your growth and without compromising levels of security or resilience. That is only achievable when you take a clear-eyed, consistent approach to your security from the start of its implementation.

The effort you make today to invest in the right infrastructure and the right integrator team determines how secure, scalable, and resilient your data centres will be tomorrow. No weak links, no compromise. Just strong foundations and the right people from day one.

Leading integrators feel like an extension of your team, safeguarding your interests and ensuring that standards never slip

Nickel-Zinc: A Powerful Data Center Alternative to Lead-Acid and Lithium-Ion

BY BRANDON SMITH, VICE PRESIDENT OF GLOBAL SALES AND PRODUCT, ZINC FIVE

NOT ALL battery chemistries are designed for the same role. Many of the energy storage technologies deployed in data centers today were developed for an earlier generation of infrastructure, defined by relatively stable loads, longer discharge durations, and different safety and sustainability expectations.

As data center power profiles evolve in response to increasingly dynamic digital workloads, higher peak power demands, and a growing focus on resilience and environmental impact, the industry is re-examining how backup power is delivered.

Nickel-zinc is a fundamentally different chemistry, purpose-built for high-power, short-duration applications where immediate power matters most, delivering rapid, repeatable response alongside superior safety and a more sustainable foundation for both emerging and traditional IT environments.

The legacy of traditional solutions

Lead-acid batteries earned their place in critical infrastructure as a predictable and cost-accessible option, particularly in facilities with minimal space constraints. Though, these systems carry large footprints, significant weight, and require frequent replacement, making them less adaptable to the changing needs of IT power infrastructure.

Lithium-ion chemistry addresses some of these challenges with higher energy density and a longer service life than lead-acid, making it an attractive upgrade where space is at a premium. However, the risk of thermal runaway

and the release of flammable, toxic gases continue to impact operator and facility safety across the data center industry.

In high-cycle, AI-driven environments, both chemistries often require system overdesign or operational compromises to manage rapid load fluctuations, temperature mitigation, and long-term infrastructure risk.

Optimizing immediate power per square foot

Nickel-zinc solutions provide up to three times the power of alternative chemistries at half the footprint of lithium-ion and one-third the weight of lead-acid. This smaller, lighter form factor simplifies installation and reduces structural requirements, allowing operators to reclaim valuable floor space and increase compute capacity without costly facility expansions.

Equally important to space utilization is matching chemistry to the application. Traditional battery energy storage systems (BESS) prioritize long-duration energy delivery and are overutilized in environments where power density and immediate

discharge capability are required. Immediate Power Solutions (IPS) are designed to deliver instantaneous, high-rate, short-duration power, and are ideal for supporting uninterruptible power supply (UPS) systems and managing AI-driven load fluctuations.

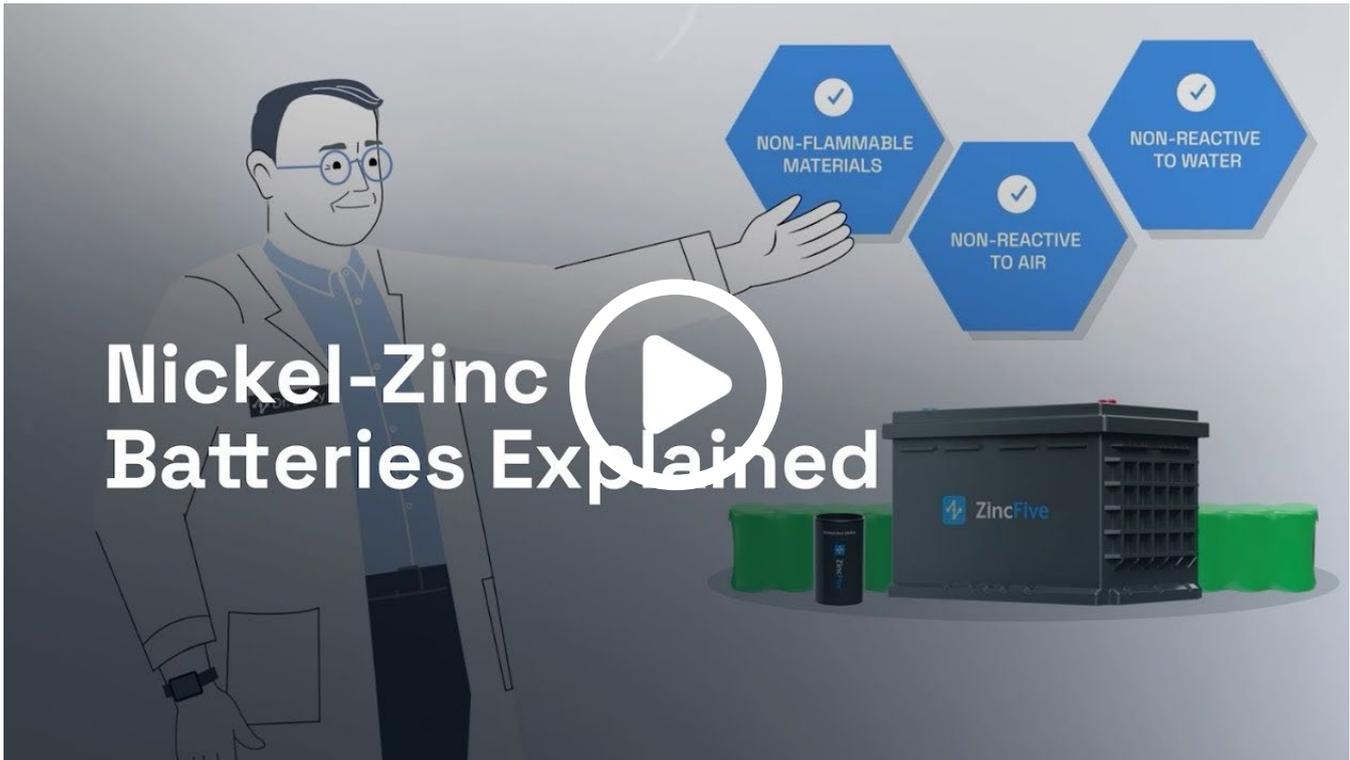
Nickel-zinc's high power density and compact footprint, alongside the safety, reliability, and sustainability benefits outlined below, make this chemistry an ideal choice for IPS applications.

Safety and reliability as a design feature

Thermal runaway events associated with lithium-ion systems have heightened safety concerns in data centers and driven the need for BMS-initiated shutdowns and extensive temperature mitigation infrastructure. Nickel-zinc technologies have inherent characteristics that minimize these challenges, capable of operating across a wider temperature range than lead-acid or lithium-ion without any risk of thermal runaway at the cell level.

Nickel-zinc systems also continue operating even if individual cells weaken or fail, supporting





Nickel-Zinc Batteries Explained

uninterrupted performance and reducing the risk of unplanned downtime.

These safety and reliability measures reduce reliance on external control mechanisms, improving operational flexibility while minimizing the likelihood of costly and dangerous failure events.

Total cost of ownership over the full lifecycle

Total cost of ownership (TCO) accounts for maintenance, replacement frequency, cooling requirements, downtime risk over the life of the facility, and post-life material management. With comparatively shorter lifespans, lead-acid systems drive higher long-term costs through frequent replacements, intensive maintenance, and inefficient space usage. Despite a longer service life, lithium-ion solutions introduce additional expenses tied to cooling systems, safety controls, and protections required to manage thermal risk.

Nickel-zinc batteries offer a service life of up to 15 years, significantly reducing replacement cycles compared to lead-acid and narrowing the lifecycle cost gap with lithium-ion. Lower maintenance demands, reduced cooling dependency, and safe operation without complex mitigation systems contribute to more predictable

operating costs over time. When combined with footprint reduction and superior safety advantages, nickel-zinc chemistry delivers a significantly lower total cost of ownership than lithium-ion or lead-acid chemistries over the full product lifespan.

Sustainability without tradeoffs

Regulatory pressure, corporate climate standards, and stakeholder expectations have elevated sustainability to a core decision criterion for infrastructure investments. Data center operators agree, as 87% of respondents from the [2025 Data Center Energy Storage Industry Insights Report](#) consider sustainability a priority, up from 81% in 2024.

Nickel-zinc batteries are built from conflict-free, widely available materials that are over 90% recyclable at end of life. With a lifespan up to 3x longer than lead-acid solutions, nickel-zinc systems reduce waste and deliver 25-50% lower greenhouse gas emissions compared to lead-acid and lithium-ion alternatives.

To support informed sustainability decisions, ZincFive's BC 2 UPS Battery Cabinet has a UL-certified Environmental Product Declaration and is included in the PEP Ecopassport database. Nickel-zinc continues to demonstrate a significantly lower climate impact than alternative

chemistries without compromising operational excellence.

A solution for AI workloads

AI workloads place unique demands on data center power systems. GPU-based processing requires power in milliseconds, often multiple times per minute, creating dynamic load spikes that can exceed 15x idle power levels and strain internal power infrastructure.

To manage these pulses, alternative chemistries often require overdesigning UPS systems at two to three times the battery quantity, along with supplemental technologies such as supercapacitors, cooling systems, or external energy storage systems dedicated to short-duration power events. These strategies increase cost, footprint, and system complexity, while still requiring graphic processing units (GPUs) to operate below peak capacity due to power delivery constraints.

Purpose-built for high-power, short-duration applications, nickel-zinc chemistry is capable of immediate response, rapid recharge, and millions of repeatable cycles over its lifespan. By absorbing AI-driven power spikes within the existing UPS footprint, nickel-zinc enables full utilization of GPU clusters without constraining performance and reduces the need for overdesign or parallel mitigation technologies.

ZincFive's BC 2 AI UPS Battery Cabinet is the dual-use nickel-zinc powered solution that allows data centers to manage AI power volatility using existing infrastructure and footprints, rather than relying on capital-intensive overbuilds.

A chemistry aligned with the future

Nickel-zinc battery technology provides backup power systems with higher power density, faster response times, uncompromising safety, and stronger sustainability standards.

Proven across millions of operating hours and engineered specifically for high-power, mission-critical environments, nickel-zinc addresses the shortcomings of alternative chemistries without introducing new compromises.

As data center power demands continue to evolve, nickel-zinc stands out as a future-proof, AI-ready solution for reliable, safe, and sustainable data center power.



ROUNDTABLE

Developing Digital Infrastructure In A Hybrid World



- Based around a hot topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion
- Moderated by editor, Phil also, this can include 3 speakers
- Questions prepared and shared in advance

Cost: £5995

Contact: Jackie Cannon jackie.cannon@angelbc.com

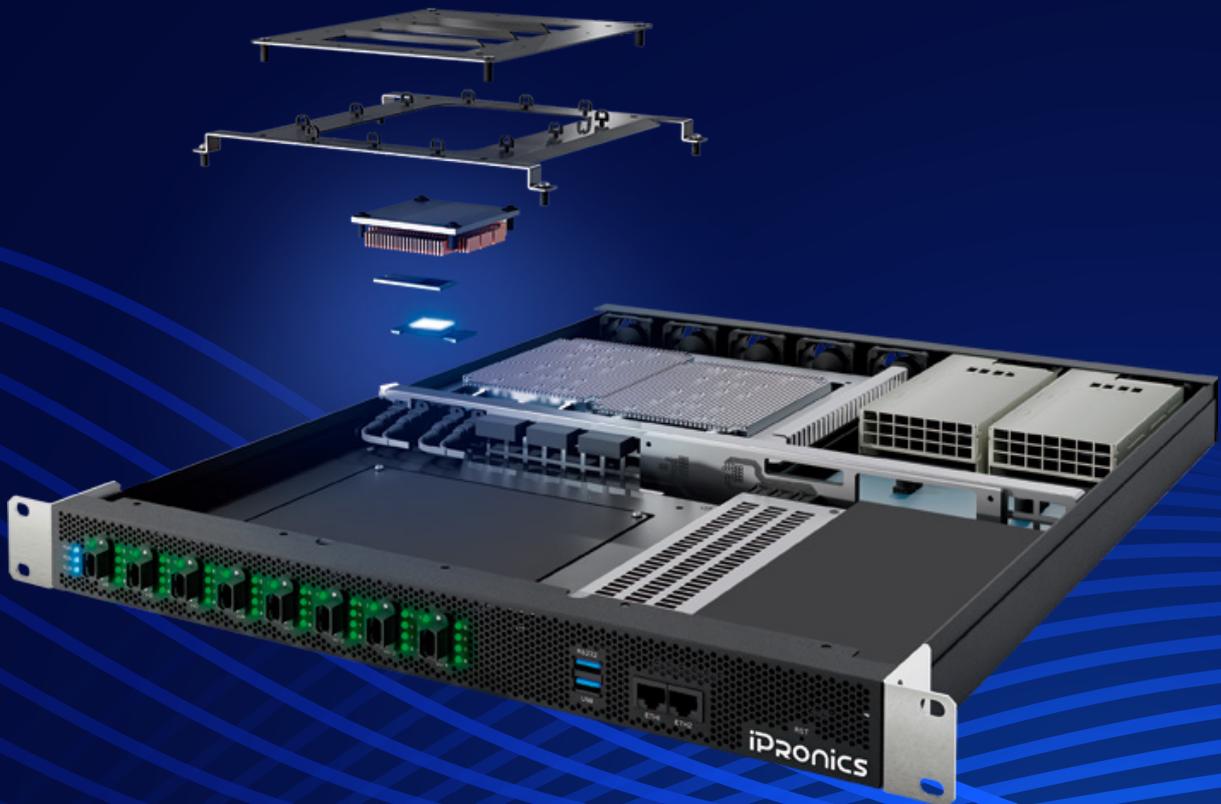
**ANGEL
EVENTS**

iPRONICS

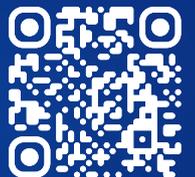
AI NETWORKING AT THE SPEED OF LIGHT

The first rack-ready optical switch built for AI. Reconfigurable in real time to maximize GPU performance, cut energy use, and scale without costly overhauls.

INNOVATION



www.ipronics.com



The top trends impacting Infrastructure and Operations for 2026

Gartner, Inc. has highlighted the top trends that will have a significant impact on infrastructure and operations (I&O) for the next 12-18 months.

“I&O leaders must be aware of all of these trends and prepare to act on the ones that are most likely to impact their organizations, so that they’ll be able to adapt, respond effectively, and drive innovation,” said Jeffrey Hewitt, VP Analyst at Gartner. “By understanding the full impact of these emerging trends, enterprises can implement effective tactics to respond, get ahead of the curve, and maximize the value of their I&O operations in 2026.”

Gartner identified six key trends impacting I&O over the next year (see Figure 1)

Trend No. 1: Hybrid computing

Hybrid computing is an emergent style that orchestrates across diverse, and sometimes incompatible, compute, storage, and network mechanisms. It enables I&O leaders to future-proof infrastructure investments with a composable and extensible compute fabric, while maximizing the value of emerging technologies by combining their strengths.

“Hybrid computing will force I&O leaders to adopt composable business and technology architecture as part of a long-term strategy for building systems and applications,” said Hewitt.

Trend No. 2: Agentic AI

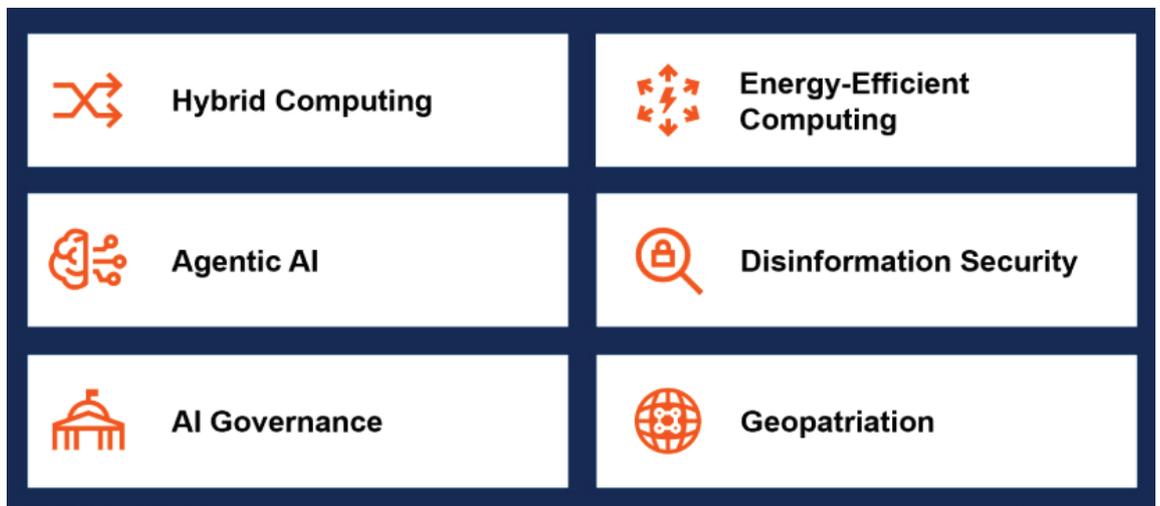
“AI is one of the top three priorities for CIO’s, and agentic AI is a very beneficial subset of that,” said Hewitt. “Agentic AI provides a significant opportunity for I&O leaders in that it enables performance gains through time savings, which will increase over time as systems evolve. It can support I&O by quickly analyzing complex datasets, identifying patterns and acting autonomously.”

Trend No. 3: AI Governance Platforms

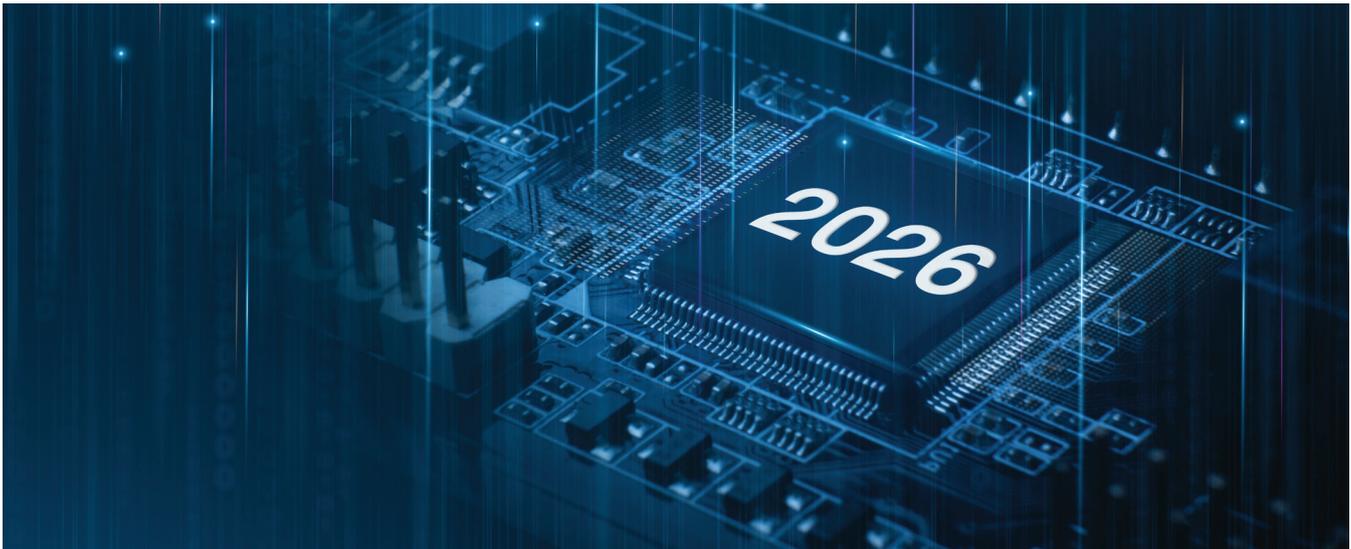
AI governance is the process of creating policies, assigning decision rights, and ensuring organizational accountability for risks and decisions related to the application and use of AI techniques. AI governance platforms oversee and manage AI systems by incorporating responsible AI practices. These platforms also address potential compliance and business risks, including bias, lack of transparency, data protection and privacy issues, model evaluation and validation, and security threats.

Trend No. 4: Energy-Efficient Computing

As a sort of subset of hybrid computing, energy-efficient computing is a package of technologies



➤ Figure 1: Gartner Top I&O Trends for 2026. I&O = Infrastructure and operations
Source: Gartner (December 2025).



and practices to reduce the energy consumption and carbon footprint of IT systems.

Energy-efficient computing enables I&O leaders to significantly impact power and environmental aspects in a [sustainable](#) way. For example, I&O leaders can create positive business outcomes by implementing tailored long-term strategies using novel and emerging technologies, such as optical computing and neuromorphic systems.

Trend No. 5: Disinformation Security

[Disinformation security](#) is a suite of technologies that can address disinformation to help enterprises discern trust, protect their brand and secure their online presence. It represents an expanding category of technologies and practices, covering deepfake detection, impersonation prevention and reputation protection.

“Given the evolving technology landscape, disinformation security will enable I&O leaders to ensure trust in communications, identity and reputation,” said Hewitt.

Trend No. 6: Geopatriation

Geopatriation is the relocation of workloads and applications from global [cloud](#) hyperscalers to regional or national alternatives due to geopolitical uncertainty.

“[Geopatriation](#) is an extension of a previous trend called ‘nationalism versus globalism,” said Hewitt. “Arguably, it goes beyond cloud from just data sovereignty to operational sovereignty to technical sovereignty. Geopatriation empowers I&O to reduce geopolitical risks and address specific sovereignty requirements. It also enables I&O leaders to support and increase the independence of domestic economies.”

AI spending to total \$2.5 trillion in 2026

Worldwide spending on AI is forecast to total \$2.52 trillion in 2026, a 44% increase year-over-year,

according to Gartner, Inc. a business and technology insights company.

“AI adoption is fundamentally shaped by the readiness of both human capital and organizational processes, not merely by financial investment,” said [John-David Lovelock](#), Distinguished VP Analyst at Gartner. “Organizations with greater experiential maturity and self-awareness are increasingly prioritizing proven outcomes over speculative potential.

“Because AI is in the [Trough of Disillusionment](#) throughout 2026, it will most often be sold to enterprises by their incumbent software provider rather than bought as part of a new moonshot project,” said Lovelock. “The improved predictability of ROI must occur before AI can truly be scaled up by the enterprise.”

[Building](#) AI foundations alone will drive a 49% increase in spending on AI-optimized servers for 2026, representing 17% of total AI spending. AI infrastructure will also add \$401 billion in spending in 2026 as a result of technology providers building out AI foundations (see Table 1).

Market	2025	2026	2027
AI Services	439,438	588,645	761,042
AI Cybersecurity	25,920	51,347	85,997
AI Software	283,136	452,458	636,146
AI Models	14,416	26,380	43,449
AI Platforms for Data Science and Machine Learning	21,868	31,120	44,482
AI Application Development Platforms	6,587	8,416	10,922
AI Data	827	3,119	6,440
AI Infrastructure	964,960	1,366,360	1,748,212
Total AI Spending	1,757,152	2,527,845	3,336,690

➤ Table 1: Worldwide AI Spending by Market, 2025-2027 (Millions of U.S. Dollars). Source: Gartner (December 2025).



The greening of the Internet Exchange Point



Following the recent COP30, the conversation around digital sustainability has never been more urgent.

BY JENNIFER HOLMES, CEO OF THE LONDON INTERNET EXCHANGE (LINX)

DESPITE TECHNOLOGY underpinning every aspect of modern life, its environmental footprint rarely features in climate discussions. The global exchange of data through AI, streaming and cloud computing is expanding rapidly, and is an invisible driver of energy demand consumption across networks and data centres.

But invisible does not mean inconsequential.

Global data creation is increasing by around 23 per cent each year, while the systems that support it already account for 1-2 per cent of global electricity use. If this digital infrastructure is to remain a force for innovation and inclusion, it must also become a model for responsible growth, reducing its carbon impact without compromising connectivity.

The carbon cost of connectivity

Every byte of data travels through a complex ecosystem of cables, routers, switches and servers. When that journey is unnecessarily long, what's known as 'tromboning', the energy cost multiplies. A video call between two people in the same city might be routed halfway across the world and back before reaching its destination with each detour consuming power and increasing latency.

As demand for low-latency services such as online gaming, video conferencing and AI powered applications grows, so too does the pressure on the network to deliver fast, stable connectivity. This demand must be met efficiently, otherwise, the very systems designed to enhance human productivity could inadvertently accelerate our carbon footprint.

The role of the internet exchange point

One of the least visible yet most vital parts of the ecosystem is the Internet Exchange Point (IXP). IXPs act as neutral meeting points where networks, such as internet service providers, content delivery networks and cloud platforms, interconnect directly. This local interconnection helps reduce the distance data must travel, cutting both latency and unnecessary energy use.

By facilitating more efficient routing, IXPs play a quiet but significant role in digital sustainability. When networks exchange traffic locally rather than sending it through multiple intermediaries or long-haul routes, energy demand is reduced.

Moreover, localised routing strengthens resilience. Regional networks can continue operating even if international

links are disrupted, reducing reliance on long-distance infrastructure.

IXPs sit at an important junction in the sustainability conversation, where operational efficiency meets measurable environmental responsibility. Organisations, for example, across the exchange ecosystem are now measuring Scope 1 and 2 emissions and, increasingly, Scope 3 emissions related to supply chains and partner operations. Understanding and tracking this impact enables more informed choices about infrastructure design, energy procurement and vendor selection.

Collaboration with data centre partners to improve energy efficiency and develop clearer ways of reporting carbon impact is gaining traction industry-wide. This move toward transparency and shared metrics represents the next stage of maturity for network sustainability, one grounded in accountability rather than aspiration.

The social governance dimension

Beyond their technical value, IXPs embody a powerful model for social governance responsibility in the digital age. Their member-driven structures ensure transparency, accountability and diversity of voices in decision making, principles that mirror the collaborative nature of the internet itself.

Many IXPs invest in education and capacity building through initiatives such as open technical training, youth engagement and digital-skills programmes. Global training initiatives have shown how knowledge sharing and skills development can

empower engineers, students and policy makers alike.

By fostering inclusion and technical competence, the IXP community contributes not only to the efficiency of global connectivity but to ensuring the benefits of a connected world are distributed more evenly. In this way, social value and sustainability are intertwined: a greener internet must also be a more equitable and transparent one. When IXPs dedicate time to social and community efforts, the impact is significant and more should be looking to make this a priority to ensure this responsibility is not shouldered by a few, but embraced collectively across the industry.

Building data 3.0: Collaboration for climate transparency

If the last decade was about connecting the world, the next must be about connecting responsibly. “Data 3.0” should not only describe an era of intelligent, AI-driven applications but also one of accountability, where every layer of the digital ecosystem, from hardware to software, is designed with sustainability in mind.

Achieving this vision requires collaboration far beyond the data-centre walls. Governments and regulators have a key role to play in setting frameworks for carbon transparency across digital infrastructure. Energy-efficient routing, renewable integration and standardised sustainability reporting should become shared priorities across the global internet community.

Encouragingly, within the connectivity eco-system there are signs of collective

progress. Industry bodies and member-driven organisations are publishing their first ESG reports, committing to annual carbon footprint disclosures and long-term reduction strategies. These steps not only improve accountability but also build a foundation for consistent environmental benchmarking across the sector.

Education and inclusion, core parts of social governance, will continue to underpin this journey. A sustainable network is, ultimately, one that empowers people as much as it transmits data.

A greener internet is within reach

The internet’s growth is both inevitable and essential. Our societies, economies and climate solutions depend on it. From precision agriculture to renewable energy forecasting, digital infrastructure is a critical enabler of decarbonisation, but it must be accountable for its own impact.

Sustainability in the digital realm is not about halting progress but about re-engineering it. From local interconnection at IXPs to transparent carbon accounting and community education, each step toward efficiency brings us closer to a greener, more resilient internet.

As COP30 convened global leaders to chart the next chapter of climate action, its time the digital world takes its place at the forefront of the conversation. The networks that connect us also connect our future and their sustainability will define not just the speed of our data but the health of our planet.



Marketing green data centres: how to turn sustainability into strategic value



Over the last decade, we've seen data centre demand grow rapidly to support the rise of Artificial Intelligence (AI), cloud and hyperscaler workloads. As awareness of environmental impact increases, expectations are rising beyond just operational efficiency.

BY MATTHEW WHALLEY, MANAGING DIRECTOR, ILEX CONTENT STRATEGIES

DATA CENTRE OPERATORS can't just innovate in Environmental, Social and Governance (ESG), they must also be transparent about their strategy, roadmap and results. This means clearly communicating how new technologies and initiatives are delivering measurable sustainability gains. By making these efforts visible and quick to understand, operators can build trust, prove value, and show how their role in environmental leadership supports long-term growth.

AI is a massive driver for rising energy demands in data centres, with electricity demands from AI-optimised data centres expected to more than quadruple by 2030, according to the International Energy Agency. Digital infrastructure providers, service providers and hyperscalers need clear, credible marketing strategies to communicate how their facilities are meeting ESG expectations.

What was once seen as a compliance issue is now a core driver of infrastructure decisions. Customers, investors and regulators are holding providers accountable for the environmental footprint of digital infrastructure. This means the way that data centre operators communicate their green credentials has become just as important as the tech powering the operations.

Sustainability under the microscope

The energy consumption of data centres is under the magnifying glass today more than ever before. As facilities scale to meet rising AI and cloud demands,

many are facing growing questions about their environmental impact. This ranges from Scope 3 emissions to water usage and e-waste.

Many data centres have made huge progress in improving their Power Usage Effectiveness (PUE) or have made switches to renewable energy. The challenge is how to communicate these advances.

Environmental performance is often only highlighted in complex, jargon-heavy technical reports. This creates a gap between what businesses are doing and what stakeholders actually understand. This 'credibility gap' risks undermining progress and reducing trust.

Sustainability expectations are also becoming a standard in tenders, procurement processes and investor

evaluations. Without clear and verifiable messaging around ESG practices, providers will lose out, no matter how sustainable operations really are.

Turning ESG into an asset

To shift sustainability from a compliance box-ticking exercise to a real competitive advantage, data centre providers must treat it as a core pillar of their brand and messaging. This means translating environmental performance into meaningful business outcomes for customers.

The most effective strategies will combine metrics with real-world impact. For example, highlighting how low-carbon operations reduce total cost of ownership for customers, or how renewable sourcing supports the clients' own ESG goals.



We need to move from abstract reporting to clear business value.

Transparency should be next on the list of priorities. ESG claims must be backed by independent verification such as International Organization for Standardization (ISO) certifications, science-based targets and Global Real Estate Sustainability Benchmark (GRESB) rankings. By being able to prove ESG claims, data centre providers can avoid accusations of greenwashing and improve their reputation. Third-party validation is becoming the language of trust in sustainability communications.

It is crucial that sustainability messaging is integrated, not isolated. Rather than siloing ESG into a dedicated report or web page, organisations should weave it through all communications, from product position and customer communications to analyst briefings and leadership content. Every touchpoint is an opportunity to build trust.

What does the buyer actually want?

To make sustainability messaging effective, data centre operators must understand what enterprise buyers and procurement teams are really looking for. Environmental goals are important, but they are often secondary to other key business priorities such as cost-efficiency, resilience, uptime and regulatory alignment.

Sustainability stories should therefore highlight how green operations help mitigate risk, future-proof operations, reduce costs over time and align with evolving compliance frameworks. Showcasing a reduction in operational energy use doesn't just signal climate responsibility. It shows price stability which is a key driver in volatile energy markets.

Offering location-specific renewable power agreements can demonstrate flexibility, helping global customers meet their own regional sustainability



mandates. Instead of abstract carbon metrics, companies should focus on the outcomes. How does sustainability make life easier, cheaper and less risky for customers? When messaging starts from a buyer-first position, sustainability becomes more than just a value. It becomes a differentiator.

Embedding ESG into our messaging

Effective ESG marketing is about who says it, where, and how often. Sustainability should be part of the brand architecture, not a peripheral message. This requires coordinated effort across marketing, sales, product, operations and leadership.

From web copy and RFP responses to social content and executive thought leadership, ESG must be consistently embedded across all customer-facing and investor-facing assets.

When done correctly, this doesn't dilute the message. It reinforces it with authority, repetition and clarity.

Internal enablement matters too. Equipping teams to confidently discuss ESG initiatives ensures the message lands consistently at every stage of the sales and partnership cycle.

Green as a long-term growth driver

The sustainability conversation in digital infrastructure is only going to increase. New regulations, rising energy costs and investor pressure will continue to push the industry towards greater environmental accountability. At the same time, customers will increasingly look for partners who can prove their credentials, not just state them.

This presents an opportunity for data centre operators and telecom providers to lead from the front. By aligning sustainability with strategic marketing, brands can set themselves apart, strengthen stakeholder relationships and future-proof their value proposition. As AI continues to scale and global infrastructure expands, the question won't just be how much energy is being used, but how responsibly it's being managed, and how credibly that message is being delivered. The companies that thrive will be those who understand that environmental performance and brand performance are now deeply connected.

Sustainability isn't just a technical challenge. It's a communications challenge. The winners will be those who can communicate their ESG story with clarity, credibility and impact.

The sustainability conversation in digital infrastructure is only going to increase. New regulations, rising energy costs and investor pressure will continue to push the industry towards greater environmental accountability

Aalberts: the lifeblood of every modern data centre

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

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

The Aalberts advantage

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

- ### • The most complete portfolio in the industry

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

- ### • End-to-end project execution

From early design and specification through prefabrication, logistics and installation support, the Aalberts focus is on taking friction and risk out of the delivery process; fewer suppliers equal fewer coordination risks, plus a faster build versus traditional build

- ### • Scalable, standardised systems

Solutions are engineered to deliver consistent performance across multisite programmes, helping to reduce operational costs, support SG targets and to maximise uptime with repeatable, factorytested designs



From plantroom specialist to fullbuilding partner

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

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

The Aalberts portfolio is the widest in the industry. Thanks to this range and technical knowledge, the available scope is broad. With plant and prefabricated elements now sitting alongside valves, connection

technologies and piping, Aalberts can support every major fluid system in the building.

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

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

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

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

End-to-end support for a changing market

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

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



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

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

Benefits that matter

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

Design and construction

- Assured compliance and safety**
 Navigating regional regulations across multiple projects can be complex. Aalberts supports project teams early in the planning phase, embedding compliance into the design through proven, factory-tested solutions and established safety standards, which reduce risk and give regulatory confidence

to all stakeholders, from Day 1. Aalberts integrated piping systems have recently upgraded and extended their valve range suit both FWS and TCS requirements, reflecting the constant focus on evolution and advancement within the business.

- Quicker, more predictable build programmes**
 Prefabricated elements enable parallel offsite construction, shortening build times and improving cost certainty compared with traditional, linear installation approaches
- Scalable, future-ready design**
 Modular platforms make it easier to replicate proven designs, expand capacity and adapt to evolving power densities or cooling technologies. Aalberts’ engineering teams support consultants with modelling and system selection to deliver flexible, site-specific solutions without sacrifice

System performance

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

- **Lower lifecycle energy and operating costs**

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

- **Reliable operation and high uptime**

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

The element of prefabrication

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

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

- **Speed without compromise**

Prefabrication within data centre builds has become a practical lever for programme certainty. The integration of pressure maintenance, expansion management and vacuum degassing into factory-built modules that arrive on site pretested and ready for connection to the wider cooling system is a straightforward method for practical savings (time and labour) on site.

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

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

- **Example: Building resilience offsite**

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

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

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

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

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

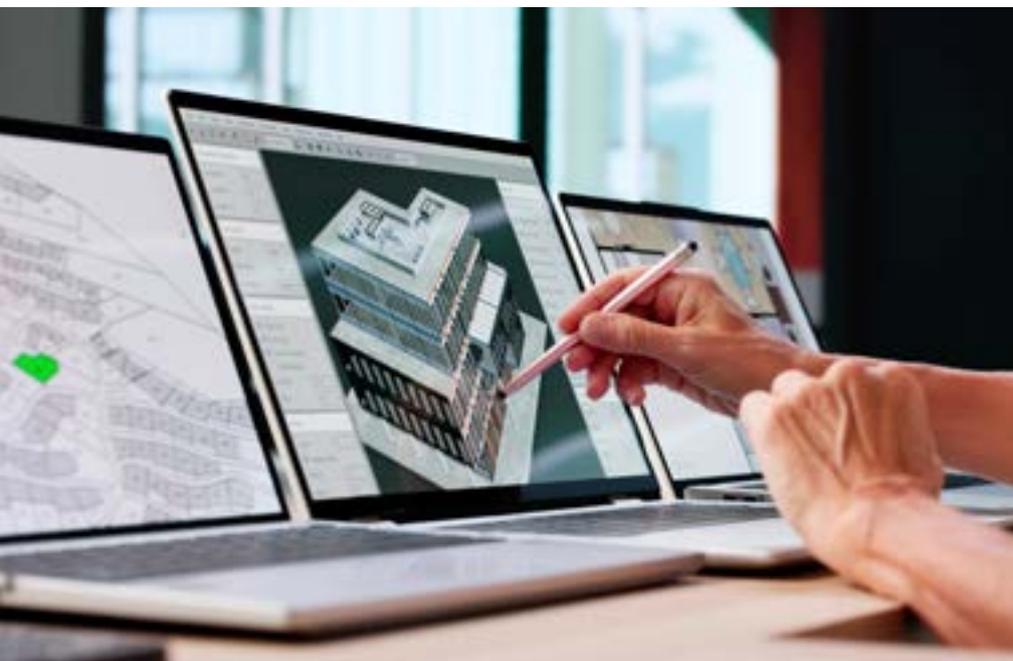
Visualising the solution: an immersive 3D experience

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

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

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

<https://tinyurl.com/aalberts3Dtour>





ZincFive

BC 2 AI

**Designed for AI.
No Tradeoffs in
Safety, Space,
or Sustainability.**

LEARN MORE

zincfive.com

Inside the AI Rack:

Power, fabric, and the optical foundations of AI infrastructure

AI infrastructure is entering a new era. Server clusters now operate as rack-scale supercomputers at power densities of 50–120 kW, with high-speed fabrics connecting dozens of accelerators per rack. This article, which supports Part 2 of AFL's three-part Hyperscale Market Shifts series, serves to highlight how AI pods scale across data halls, how the optical layer supports operational reliability, and how advanced fiber architecture supports long-term growth.

Rack architecture and power delivery

Advancing from nodes to rack-scale systems

The basic unit of AI compute is changing. Where clusters once scaled by adding servers, today's AI workloads are planned in racks and pods. A single rack can integrate dozens of accelerators joined by a high-bandwidth scale-up fabric, behaving as one large processor. Capacity planning, resiliency, and expansion strategies increasingly revolve around rack-scale builds rather than individual nodes.

Power, cooling, and rack density constraints

As rack power pushes beyond 100 kW, air cooling alone is no longer sufficient. Direct-to-chip liquid cooling and warm-water loops are becoming standard, influencing both rack design and facility infrastructure. However, rear-of-rack space is already crowded, leaving little margin for ad-hoc cabling. Therefore, coordinated system design must now include power, cooling, and connectivity considerations from the beginning.

Scale-up and scale-out fabrics

Inside the rack, scale-up networks connect accelerators with extremely high bandwidth and low

latency, pushing the limits of copper interconnects. Beyond the rack, scale-out fabrics link racks into pods using Ethernet or InfiniBand at 400G and 800G, with 1.6T in development. At these speeds, fiber counts rise rapidly. A single switch face can terminate more than a thousand fiber cores, and pod-scale fabrics quickly grow to hundreds of thousands of fiber terminations.

Optical fabrics and structured cabling

Fiber Density and Physical Structure

Disciplined design is essential for both physical topology (paths, trays, routing) and logical topology (device connections, traffic flow). High-fiber-count trunks, defined distribution points, and front-of-rack access can help to streamline deployment, reduce congestion, simplify testing, and create clear paths for future upgrades.

Optical fabrics inside the rack

Co-packaged optics (CPO) and linear pluggable optics (LPO) extend reach, reduce power per bit, and allow more flexible layouts. As a result, traditional boundaries between intra- and inter-rack fabrics become less distinct, while optical layers continue to anchor long-term performance.

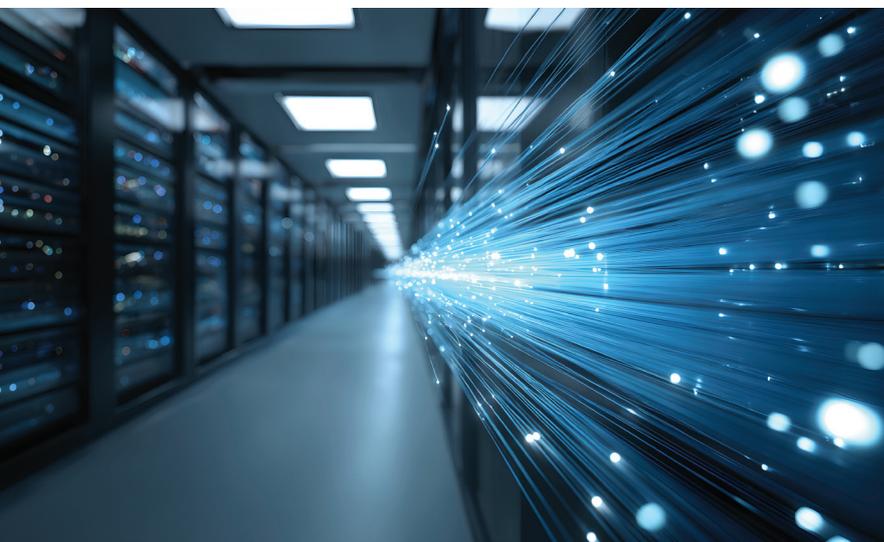
Preparing the fiber plant for future generations

To accommodate 800G today and 1.6T tomorrow, today's data center designs must integrate scalable pathways, excess tray capacity, and modular termination architectures. Using bend-insensitive single-mode fiber, high-fiber-count trunks, and standardized polarity schemes can reduce friction during upgrades, while structured testing, inspection, and documentation practices can help to preserve stability and reliability across multiple generations of optics.

Operational management and reliability

Managing Complexity at Pod Scale

Scaling racks into pods and pods into halls introduces significant operational complexity, with



a single AI pod containing hundreds of racks, thousands of switches, and hundreds of thousands of fiber terminations. Disciplined physical design, structured cabling, and thorough documentation can help technicians mitigate this complexity while maintaining operational reliability.

Reliability and failure domains

Loss of a single rack can remove dozens of accelerators, creating a substantial failure domain. Redundant fabrics, multiple power feeds, and diverse fiber routes are required to prevent localized issues from cascading across pods.

For enhanced resilience, data center planners may consider separating redundant paths, documenting interdependencies, and avoiding single points of failure in patch panels.

Key takeaways

- Rack-scale designs serve as the atomic unit of AI compute, requiring integrated planning for power, cooling, and connectivity across high-density racks and pod-scale deployments.
- Optical layers provide long-term infrastructure stability, supporting upgrades from 800G to 1.6T and beyond while accommodating evolving hardware generations and increasing bandwidth demands.

- Structured cabling and modular fiber pathways reduce congestion, simplify testing, and enable predictable upgrades across pods, rows, and complex hyperscale environments.
- Disciplined operational practices and thorough documentation mitigate risk, improve scalability, and sustain availability in high-density, high-power AI data center deployments.
- Redundant fabrics, multiple power feeds, and diverse fiber routes enhance resilience and limit failure domains while maintaining high throughput and minimizing downtime.

Explore more expert insights

For in-depth guidance on hyperscale infrastructure, fiber management, and AI data center design, access AFL's full library of blogs, white papers, and technical resources: [AFL Insights](#)

Access the complete Hyperscale Market Shifts series here:

Part 1 – AI, Neoclouds, and the New Limits of Data Centers

Part 3 – Advanced Technologies and Macro Trends Affecting AI Deployments

DCS ROUNDTABLE
Developing Digital Infrastructure In A Hybrid World

- Based around a hot topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion
- Moderated by editor, Phil alsop, this can include 3 speakers
- Questions prepared and shared in advance

Cost: £5995

Contact: Jackie Cannon jackie.cannon@angelbc.com

ANGEL EVENTS

iPronics ONE – The optical engine redefining AI networking

As artificial intelligence reshapes digital infrastructure at unprecedented speed, data centers face a critical reality: traditional electronic networking is no longer enough.

AI WORKLOADS demand massive bandwidth, instantaneous responsiveness, flexible topologies, and near-zero downtime. Yet electronic packet switching (EPS) – the decades-old foundation of today’s networks – was designed for a different era, one defined by modest, unpredictable data flows rather than synchronized GPU-to-GPU communication at scale.

iPronics ONE, the Optical Networking Engine, redefines what AI infrastructure can be. Built on programmable silicon photonics, it delivers ultrafast, lossless optical switching directly at the physical layer – unlocking a new generation of

high-performance, energy-efficient, and topology-adaptive AI clusters.

A new paradigm: faster, simpler, optical

AI training and inference require constant data exchange among GPUs, often at volumes that grow exponentially with model size. This communication must be deterministic, synchronized, and extremely fast. Yet EPS architectures impose latency, jitter, high power consumption, and the need for multiple electro-optical conversions, adding complexity and cost at every hop.

iPronics ONE eliminates these bottlenecks with a solid-state,

fast-reconfigurable, lossless Optical Circuit Switch (OCS) based on fully programmable silicon photonics. Its sub-millisecond reconfiguration enables real-time adaptation to the communication patterns of each AI workload. This is AI networking at the speed of light, engineered for the performance and scale requirements of the world’s most advanced compute clusters.

Built for modern AI data centers

iPronics ONE integrates seamlessly into AI clusters and next-generation data centers, introducing capabilities that fundamentally transform the physical layer:



- Onthefly topology reconfiguration**
AI workloads evolve through multiple communication phases. With iPronics ONE, the network can instantly reshape itself to match those phases, maximizing cluster efficiency.
- Full Oband operation**
Compatibility with DR, FR, and other standard optical modules ensures smooth integration into existing infrastructure and readiness for future transceiver formats.
- 32radix switching (scalable to 64, 128, 256)**
Expansion is built into the architecture, enabling growth without redesign.
- Integrated telemetry & perchannel gain control**
Realtime monitoring and optical gain ensure stable, lossless connectivity across every link – essential for synchronized GPU performance.
- Submillisecond reconfiguration**
The switch responds faster than workload transitions, ensuring the network never becomes the bottleneck.
- SDNcompatible API**
iPronics ONE integrates natively with softwaredefined data centers, allowing orchestration platforms to control the optical layer and automate topology changes.

The result: a new category of networking hardware that combines the precision of silicon photonics with the flexibility of cloudscale programmability.

Why optical switching is now essential

As AI platforms scale, EPS alone cannot meet the combined demands of performance, efficiency, and cost. iPronics ONE addresses these challenges head-on, delivering advantages that directly impact operations:

- Enhanced AI performance**
Lossless connections, ultralow latency, and high bandwidth enable faster, more predictable training cycles, improving throughput and shortening development timelines.
- Improved resource utilization**
Machinelearning traffic is inherently predictable. Dynamic topology adaptation ensures GPUs – and other highvalue compute resources – operate at peak efficiency.
- Lower Total Cost of Ownership (TCO)**

iPronics ONE reduces longterm costs by offering:

- Significantly lower power consumption
- Up to 50% fewer required transceivers
- Simplified network management through optical programmability

These factors collectively deliver a strong return on infrastructure investments.

- Greater agility and scalability**
AI workloads evolve rapidly. With submillisecond reconfiguration, data centers can adapt instantly and scale seamlessly as demand grows.
- Higher reliability and resilience**
Automatic, ultrafast failure recovery ensures uninterrupted GPU synchronization – preventing costly disruptions in training processes.

A leap beyond traditional OCS technologies

While MEMS and LCOS optical switches served earlier networking needs, they cannot keep pace with modern AI requirements. iPronics ONE offers clear advantages:

- Solidstate switching (no moving parts)**
Standard proven CMOS manufacturing process with higher reliability and lower cost per port.
- Lossless optical paths**
Perchannel gain control with integrated SOA channel gain control eliminates the 2–3 dB losses typical of MEMS/LCOS.
- Up to 1000× faster reconfiguration**
Essential for dynamic AI workloads that require instant adaptation.
- Lower crosstalk**
Provides clean, deterministic optical signals for highprecision GPU communication.
- Bandwidthagnostic operation**
Switch ports remain data rate and protocol agnostic, compatible as networks evolve toward 800G, 1.6T, and beyond – futureproofing infrastructure.
- Automatic linkfailure recovery**
Builtin integrated telemetry provides network resilience, ensuring continuity for in resilience ensures continuity for highperformance computeperformance compute workloads.

These advantages position iPronics ONE as the most advanced optical

circuit switching solution available today.

Enabling new AI and data center architectures

iPronics ONE supports transformative use cases across both scaleup and scaleout architectures:

- Higher GPU density per rack**
OCS enables denser GPU configurations without overwhelming EPS layers.
- Dynamic topologies based on workload phase**
A level of adaptability impossible with static electronic networks.
- Fast failure recovery for uninterrupted GPU synchronization**
Essential for long, resourceintensive AI training jobs.
- Transparent evolution to higher bandwidths**
OCS platform remains the same as data rates increase – no need to upgrade switch ports.
- Simplified infrastructure and reduced complexity**
Saving 50% number of transceivers lowers cost, power use, and management overhead.

iPronics ONE is not just a switch – it is the foundation for the next generation of AI data center designs.

A new standard for AI networking

The rapid acceleration of AI is pushing traditional networking technologies to their limits. iPronics ONE introduces a new foundation: fast, programmable, lossless, and built for the speed and scale of modern AI.

With siliconphotonicsbased optical circuit switching, iPronics delivers a network platform that is:

- More efficient
- More scalable
- More resilient
- More reliable
- More cost-effective
- More futureproof
- And dramatically faster

In a world where milliseconds matter and every watt counts, iPronics ONE is the optical engine ensuring AI infrastructure keeps pace with innovation.

AI runs on computation.

The future of that computation will run on light.



DCA INSIGHTS

Steve Hone
CEO, The DCA
(Data Centre Alliance)



A belated Happy New Year

THE DCA wrapped up 2025 in style by hosting our final 10X10 Data Centre Update at PowerEx Live, London. As always, it was a buzzing event, featuring some of the best networking and social opportunities of the year. A huge thank you to our 10X10 presenters from Evapco, Legrand, RED Engineering, CMP Products, Wesco-Anixter, Sudlows, MEUC, Piller and Fresco Clean Tech Solutions, for sharing their insights and expertise. I would also like to thank the PowerEx team for their continued help and support throughout 2025!

It's now full steam ahead at DCA Towers with preparations for another busy year. The number of events and conferences across the data centre industry has grown exponentially, and The DCA will be actively attending and supporting many of these—representing both DCA Partners and the wider sector.

We are particularly looking forward to Data Centre World London 2026, where The DCA will be chairing the Keynote Theatre, hosting panel sessions, and judging the DCW Awards. DCA Partners are encouraged to visit The DCA Stand G10 to catch up with the team and take part in video interviews with Data Centre Solutions.

In 2026 we will be hosting more of the well-established 10X10 Data Centre Updates and Data Centre Transformation Conference. For full

It's now full steam ahead at DCA Towers with preparations for another busy year. The number of events and conferences across the data centre industry has grown exponentially, and The DCA will be actively attending and supporting many of these

details, please visit www.dcauk.org or follow The DCA on LinkedIn.

DATA CENTRE SOLUTIONS – ISSUE 01 2026

This DCA feature is comprised of articles from DCA Partners and industry experts. Thank you to all the authors for their insightful contributions.

- **Building Resilient Data Centre Security Operations: A Ground-UP Approach** - David Jackson, Operations Director, ICTS Ireland provides us with an interesting insight into physical security and the four pillars model used at ICTS.
- **Lifecycle Extension as a Resilience Strategy** - With AI demand straining

IT supply chains, extending the life of existing infrastructure isn't just cost control, it's a resilience strategy. *Alastair Winner, Head of Innovation and Services, Techbuyer* provides us with his thoughts.

- **All Suppliers Must Contribute to Data Centre Sustainability** - Hans Obermillacher, Business Development – Panduit. Hans explains why sustainability has become increasingly important for data centres, in their build, fit out and operation.
- **The Hidden Energy Asset Inside Every Data Centre**, By Andy Lewis, General Manager, Data Centres at Weatherite. Andy talks us through the renewed interest in heat recovery not simply as a sustainability initiative, but as a practical strategy for improving overall energy efficiency and long-term operational resilience.
- **Designed to be Built - How Modern Construction Methods Improve Cost, Programme and Quality** Zac Potts - Sudlows details how the right approach to DC design is one of the most powerful tools in risk mitigation and speed to market.

If you'd like to find out more about The DCA and how we support the sector and those working in it, drop me an email: steveh@dcauk.org

Best regards,
Steve

KEY DC INDUSTRY EVENTS Q1 2026

- **29 January** - Safety, Security & Contamination Control for the AI Age hosted by NDCA (National Data Centre Academy), Royal Leamington Spa
- **3 - 4 February** – Kickstart Europe 2026, Amsterdam RAI
- **25 February** - Platform UK 2026 hosted by Platform Markets Group at Convene, London
- **4 - 5 March** - Data Centre World (Tech Show London) at the ExCel London

Lifecycle extension as a resilience strategy



With AI demand straining IT supply chains, extending the life of existing infrastructure isn't just cost control, it's a resilience strategy.

BY ALASTAIR WINNER, HEAD OF INNOVATION AND SERVICES, TECHBUYER

The Supply Squeeze

THE AI BOOM is consuming semiconductor capacity, GPU allocation and enterprise server components at an unprecedented rate. Whether or not you're deploying AI workloads, you're now competing for silicon with organisations that are. The hyperscalers are absorbing fab capacity, and that ripples through to lead times on perfectly ordinary enterprise servers.

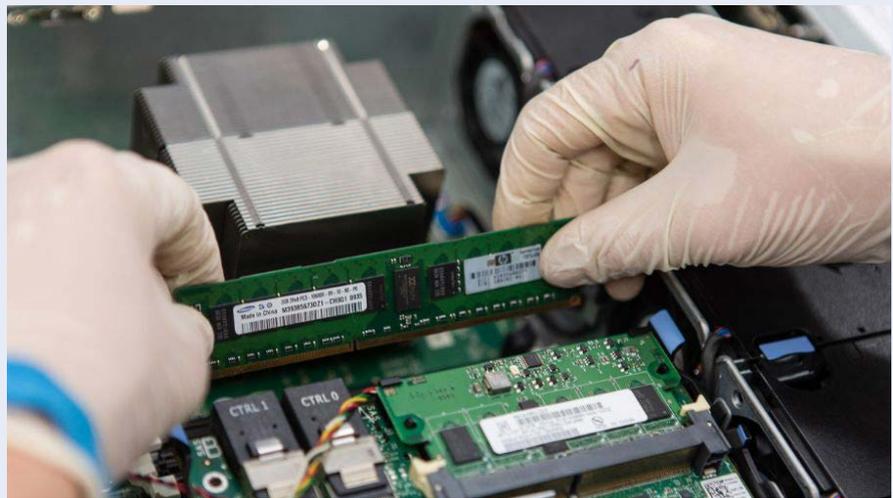
The effects are visible across the entire IT stack. Memory prices are rising as manufacturers reallocate wafer capacity from DRAM to high-margin HBM for AI infrastructure - DDR4 prices up 38–43%, persistent DDR5 shortages, and NAND contract prices climbing 5–10% per quarter. Demand is outstripping supply.

The data centre fabric faces similar pressures. Lead times for large power transformers now average 128 weeks - that's two and a half years. Liquid cooling demand is projected to grow over 1,000% by 2030, but manufacturing capacity isn't keeping pace. A survey of 280 data centre professionals found that 83% don't believe supply chains can deliver the advanced cooling technology required for AI workloads.

When equipment does become available, there's a strong incentive to over-provision. Operators burned by long lead times are pre-purchasing years in advance, forming consortiums for shared manufacturing slots, and ordering capacity they may not need for years. This hoarding behaviour tightens supply further and ties up capital in assets that may become stranded if demand projections don't materialise.

The AI reality check

Here's the uncomfortable truth: 95% of enterprise AI initiatives are failing



to deliver expected ROI, according to MIT research. Adoption is growing with 78% of organisations now using AI in at least one function, up from 55% a year earlier, but adoption doesn't equal value capture. Gartner predicted 30% of generative AI projects will be abandoned after proof of concept by the end of 2025 due to poor data quality, escalating costs or unclear business value. My sense is the number will be higher.

When we talk to IT professionals about their actual compute requirements, and not their board's AI ambitions, the vast majority are running traditional enterprise workloads. ERP systems, virtualised environments, storage. That's not going away any time soon.

The question isn't whether AI matters. It's whether betting your infrastructure strategy on AI demand that hasn't yet materialised is a sensible approach to resilience.

Extending life, reducing risk

Replacing systems wholesale and migrating to new infrastructure carries significant risk to operations, continuity

and uptime. Add extended lead times and component shortages into the mix and that risk compounds. Your N+1 architecture assumes you can source a replacement when you need it. What's your plan when lead times exceed your risk tolerance?

Extending server lifecycles through quality refurbishment and third-party maintenance is the alternative. It's not about avoiding investment it's about investing wisely. Keep proven infrastructure running reliably while freeing up resources for targeted AI experimentation where it makes sense.

Most AI inference can run on standard enterprise compute (it's model training that demands specialist hardware). The UK's energy and infrastructure constraints mean we're better positioned for inference than training anyway. For many organisations, existing equipment can support both traditional workloads and early AI use cases without a wholesale refresh.

The trust question

The obvious concern is whether the use of refurbished hardware represents a

risk for mission-critical environments? It depends entirely on the source. Enterprise hardware from certified providers goes through rigorous testing, often more thorough than factory QA, because we're validating proven components rather than catching manufacturing defects. A certified refurbished server with documented history is lower risk than grey market 'new' hardware with no traceability.

Third-party maintenance specialists often have deeper accumulated knowledge of specific platforms than original equipment manufacturers support teams. The key is access to a parts inventory, which brings us back to supply chain diversification. The secondary market represents a parallel supply chain, less exposed to the fab capacity constraints affecting new equipment.

Resilience through pragmatism
Circular economy isn't just good ethics, it's practical resilience. Extending hardware life reduces manufacturing



demand, which eases supply chain pressure for everyone. The embedded carbon in existing hardware is already spent; getting another three to five years from a server is almost always better for emissions than manufacturing a replacement.

The organisations that will thrive aren't those chasing every new trend.

They're the ones with diversified supply chains, realistic assessments of their actual workload requirements, and the discipline to sweat their assets intelligently while the market finds its equilibrium.

Resilience sometimes means not being first. It means being confident in what you're running.

THE DCA (DATA CENTRE ALLIANCE) IS THE UK TRADE ASSOCIATION FOR THE DATA CENTRE SECTOR

WE SUPPORT DATA CENTRE OWNERS & OPERATORS, SUPPLIERS ALONG WITH ALL THOSE WORKING IN THE SECTOR



We ensure the DC community stays informed and connected through our events, roundtables, articles and more.

Contact The DCA today to find out more about the benefits of our Partnerships.
T: +44 (0)845 873 4587 • E: info@dcauk.org • W: www.dcauk.org

Building resilient data centre security operations: A ground-up approach



As data centre infrastructure continues to expand across Ireland and the wider UK and European markets, the role of physical security has evolved significantly.

BY DAVID JACKSON, OPERATIONS DIRECTOR ICTS IRELAND, ICTS UK & IRELAND

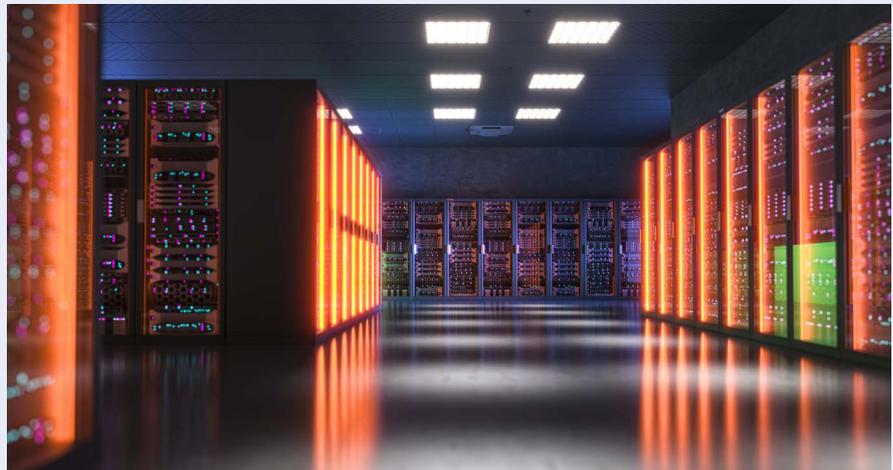
FACILITIES are no longer simply buildings requiring perimeter control. They are mission-critical operational environments supporting national infrastructure, cloud platforms, financial services and global enterprise workloads.

Within this context, security provision must move beyond static guarding models. It must integrate operational discipline, technical awareness, regulatory compliance and continuous performance management. At ICTS, our approach to data centre security is built from the ground up, focusing on people, process and operational resilience from the first day of mobilisation.

Establishing operational foundations during contract mobilisation

When ICTS assumes responsibility for a data centre contract, the initial focus is placed on establishing stable operational foundations. While all security officers are required to hold appropriate regulatory licensing through the PSA or SIA frameworks, this alone does not prepare teams for the specialist requirements of live data centre environments.

Each facility operates with its own access control architecture, zoning arrangements, client protocols, resilience requirements, and escalation structures. For this reason, ICTS implements site-specific induction programmes and data centre focused training modules during mobilisation. These programmes ensure officers understand both the physical security requirements and the operational sensitivities of working within critical national infrastructure environments.



In many transition projects, we observe that legacy service models often lack structured training depth or operational standardisation. Addressing these gaps early is essential to maintaining service continuity and reducing risk during the early contract phase.

A structured operating framework: The four pillars model

ICTS structures data centre service delivery around four operational pillars:

- Operations
- Training
- Compliance
- Customer Satisfaction

This framework provides a consistent management structure across multiple sites and allows performance to be measured, reviewed and continuously improved.

Operations: Building interoperability and service continuity

Operational resilience is a core requirement within data centre

environments. ICTS places strong emphasis on security personnel interoperability, ensuring officers are capable of performing multiple operational functions where appropriate. This approach supports business continuity during shift changes, peak demand periods, or unplanned absences.

Operational planning also incorporates structured supervision models. Frontline supervisors and area management teams maintain constant visibility of site performance, escalation activity and staff coverage. This layered management structure enables rapid decision-making and maintains consistent operational standards.

ICTS also invests in junior leadership development programmes. These initiatives focus on developing supervisory capability, communication skills, professional conduct, and operational decision-making. Strong frontline leadership directly contributes to stability, staff retention, and service quality.

Training: Developing capability for high-pressure environments

Training is treated as an operational control rather than a compliance exercise. ICTS utilises scenario-based training methodologies designed to replicate real-world data centre security incidents, including access control breaches, alarm activations, emergency response coordination, and procedural non-compliance events.

This practical training model allows officers to develop situational awareness and procedural confidence in controlled environments. ICTS applies a structured learning methodology built around demonstration, guided practice, supervised application, and operational confirmation.

In addition, ICTS incorporates the 'T-CUP' decision-making framework to support effective performance under pressure. This approach focuses on structured thinking, controlled response, and disciplined escalation pathways. When officers respond based on training rather than instinct, operational consistency improves and error rates reduce.

Ongoing validation is carried out by site supervisors and management teams through daily briefings, spot audits, and operational observations.

Compliance: maintaining standards and reducing operational risk

Data centre security environments operate within strict procedural and regulatory frameworks. ICTS ensures that all officers operate in alignment with site assignment instructions, standard operating procedures, access protocols and client compliance requirements.

Training records, certifications, and competency matrices are actively managed and reviewed. This provides management teams with visibility over capability levels and enables proactive re-training where required.

Health and safety compliance is embedded across all operations. ICTS promotes a reporting culture that encourages hazard identification, near-miss reporting, and continuous improvement. This approach protects staff welfare while maintaining safe working conditions across complex operational environments.

In many transition projects, we observe that legacy service models often lack structured training depth or operational standardisation. Addressing these gaps early is essential to maintaining service continuity and reducing risk during the early contract phase

Customer satisfaction: Delivering through partnership models

ICTS positions customer satisfaction as a performance outcome rather than a contractual obligation. Structured engagement frameworks are implemented with clients, including KPI reporting, monthly performance reviews and continuous improvement planning. These reviews provide measurable insight into service delivery performance, response times, compliance status, training progression, and risk mitigation activity. Transparency and structured communication allow both client and provider to align operational priorities and adapt service models as site requirements evolve.

By adopting a partnership-driven approach, ICTS is able to align security operations with wider client business objectives and infrastructure strategies.

Understanding the operational reality of data centre security

Physical security within data centre environments directly supports digital risk management. Officers are responsible for enforcing access control protocols, preventing unauthorised devices and storage media from entering restricted zones, and ensuring strict compliance with site rules.

Access management, screening procedures, audit logging, and controlled movement of personnel form the backbone of daily operations.

ICTS ensures that officers apply these controls consistently and without exception, regardless of operational pressure or workload volume.

Alarm monitoring and response coordination are equally critical. Teams operate against defined service level agreements and escalation frameworks to ensure rapid response and effective incident management.

Accountability across the management structure

ICTS operates on the principle that accountability must exist at every organisational level. Site supervisors, area managers, and contract leadership teams share responsibility for operational performance, training oversight, and compliance assurance. Regular performance reviews, site audits, and operational governance meetings ensure standards remain consistent and that corrective actions are implemented quickly when required.

Building professional confidence through capability development

Sustainable security performance is achieved by building confident, capable teams. ICTS focuses on developing officers who understand their responsibilities, trust their training, and operate with professionalism and consistency. When teams are properly supported and managed, they deliver stronger client engagement, improved operational outcomes, and enhanced resilience across critical infrastructure environments.

Conclusion

As data centre infrastructure continues to expand in scale and strategic importance, security delivery models must evolve accordingly. ICTS provides data centre security services built on structured leadership, operational discipline, specialist training, and partnership-driven client engagement. By establishing strong foundations at mobilisation and maintaining continuous performance management throughout the contract lifecycle, ICTS supports the secure operation of critical digital infrastructure across Ireland, the UK and beyond.

As a Group, ICTS Europe secures over 245 data centre locations across 14 countries.

All suppliers must contribute to data centre sustainability



Sustainability has become increasingly important for data centres, in their build, fit out and operation. With major UK investment commitment, in 2026, of around \$18.24bn, according to Ceres Property, the UK data centre industry has expanded its sustainability pledge through a blend of regulatory pressure, voluntary targets and innovations.

BY HANS OBERMILLACHER, BUSINESS DEVELOPMENT - PANDUIT

ACROSS EUROPE, global infrastructure providers, such as Panduit, continue to re-engineer their systems and products to align with customers and regulatory demands, including high density cabling, containment and thermal management systems targeted to data centre use cases from edge and hyperscale to colocation environments.

Sustainability has become a strategic approach and emphasises scalable physical infrastructure that can be adapted to evolving compute demands, high performance networks and greater operational density.

European data centre operators typically work within stringent standards for network performance, thermal efficiency, space utilisation and sustainability compliance. Panduit's solutions are engineered to support structured cabling, framed containment architectures and flexible pathways that enhance reliability, simplify deployment, and ensure protection for the live equipment (severs/switches, etc), and

fibre and high-performance copper cabling to ensure uptime is maintained to the highest obligations.

Image: Containment Solution

Compute and cabling containment requirements in data centres

Modern data centre facilities must seamlessly integrate compute hardware, power delivery, network cabling and containment in a unified infrastructure plan. The primary technical drivers include increased Bandwidth (800G/1.6T), reduced latency (InfiniBand for AI with 1–2 μ s), maintaining cable bend radius, protecting transmission media, and managing airflow. Containment systems are deployed to segregate hot and cold aisles so that cooling systems operate at maximum efficiency.

They also support modular scalability that allows for incremental growth without excessive restructuring of existing infrastructure, such as structured cabling.

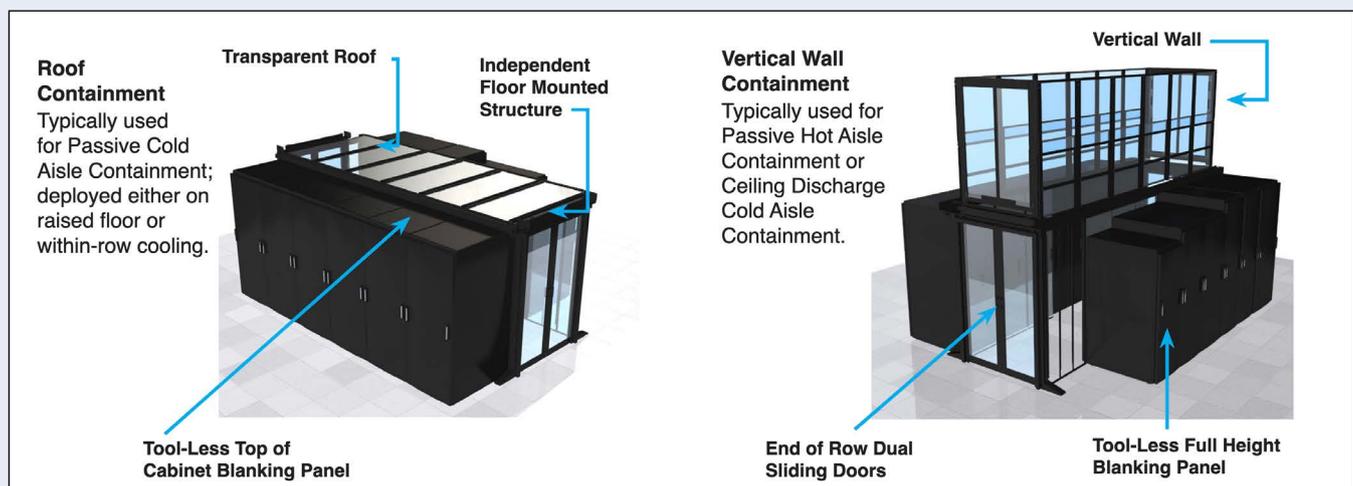
Data centre physical infrastructure can be divided into three architectural layers:

● Compute Frame Layer

This includes rack enclosures that house servers, storage, networking switches, and related hardware. These enclosures must support high-density equipment with adequate structural capacity and must be compatible with containment elements, which now include 52U server cabinet, this level of consolidated compute will often require liquid cooling solutions.

● Cabling Pathway Layer

Cable management and pathways ensure that fibre and copper cabling remains organised, protected, and accessible. Pathways must protect against signal degradation, maintain bend radii for data performance and support high volumes of cable without obstruction. Correct selection of cable can reduce cable bundle size and allow for increased cable numbers within specific cable



runs, while maintaining air circulation maintaining performance radii, and reducing cable over-heating.

○ Containment and Thermal Layer

Containment systems manage airflow and cooling efficiency. They separate cold supply air from hot exhaust air which increases cooling effectiveness, reduces wasted energy and eliminates hot air to the server air intake.

At present air cooling remains the industry's dominant cooling type, while the increasing demands from AI and HPC are driving higher power densities in data centres. These applications require liquid cooling solutions to improve energy efficiency and heat management, for the latest server pods. Data centres in advanced economies are driving over 20 percent of electricity demand growth. Liquid cooling assists sustainability as a high efficiency greener alternative to air cooling solutions.

Cabinets, cable management and containment solutions

Panduit's approach to cabinets and containment emphasises structural integrity, thermal performance and ease of installation. Cabinets are central to data centre physical infrastructure because they provide the mechanical platform for compute hardware and the structural support for cable routing and containment.

Panduit supports a variety of cabinet configurations suited to different density and deployment scenarios. Cabinets are engineered to optimise airflow around equipment and integrate with containment systems that channel intake and exhaust air to reduce cooling loads. Their design accommodates varying equipment heights, cable entry points and load capacities, as well as liquid cooling designs.

Cable management and routing systems are engineered to protect and organise both copper and fibre cabling. Structured cabling standards require a controlled infrastructure that maintains minimum bend radii, prevents crushing and simplifies access for moves adds and changes (MACs). Systems such as wire baskets and raceways (LSZH. Low Smoke Zero Halogen, according to IEC 60332-1-2) provide overhead or underfloor cable support to separate

cabling from other services, ensure mechanical protection and improve space utilisation.

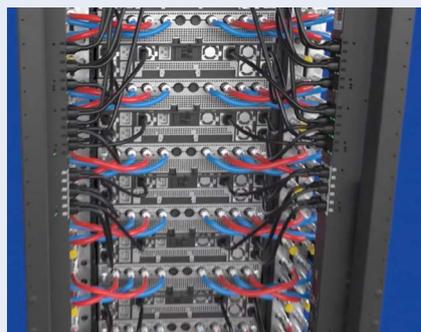
Thermal management and aisle containment are core to sustainable data centre operations. Effective containment reduces recirculation of hot exhaust air into cold supply air and maintains defined thermal zones that help improve cooling efficiency. Modular containment systems can be retrofitted into existing facilities or integrated into greenfield builds to balance temperature control with physical access.

Pathway solutions such as overhead wire basket routing and structured raceways maintain cable integrity throughout the facility. They provide a continuous and organised cable path that protects against mechanical damage, supports high cable densities and maintains consistent performance over the cable lifecycle.

Manufacturing and supply chain in Europe

Panduit supports regional supply chain and manufacturing capabilities that align with European demand patterns. This includes preconfigured infrastructure modules, local stockholding and supply interfaces with tier one distributors and system integrators.

Panduit's simplified order-to-installation model enables data centre operators to order cabinets and containment systems that have integrated Infrastructure components from a single SKU. This approach reduces packaging waste, vehicle journeys and installation time while improving predictability of deployment. Preconfigured systems can be deployed with tested quality and ensure compatibility across



➤ Panduit FlexFusion Direct-to-Chip Liquid Cooling Cabinet.

containment and cabling layers. The reduction in packaging and consolidated shipments align with sustainability and logistics optimisation objectives that are central to European data centre operations.

Sustainability: Infrastructure, packaging and supply chains

Sustainability strategies are increasingly fundamental to data centre design and operation. Panduit's sustainability focus addresses lifecycle impacts from raw materials through installation and end-of-life disposal. These policies encompass environmental stewardship, supplier codes of conduct and waste reduction initiatives. Panduit promotes engineering solutions that reduce material use and energy consumption. Smaller diameter cabling systems require less material while meeting performance criteria. These designs also improve shipping efficiency by allowing more product to be transported per pallet which reduces transport emissions.

Packaging initiatives

Packaging changes have removed single-use plastics for patch cords and transitioned to Forest Stewardship Council-certified paper labelling. This shift eliminates millions of plastic bags annually and reduces tape and excess materials in packaging. Panduit also operates a Packaging Council to track trends, increase recyclability and identify reusable packaging options. Installation documentation is increasingly provided digitally to reduce paper use.

Panduit requires suppliers to adhere to environmental standards and codes of conduct that include labour, health, safety and environmental criteria. Environmental Product Declarations (EPDs) measure impacts of material choices and contribute to customer sustainability goals like LEED or BREEAM points. Within data centre infrastructure design, thermal management and containment systems contribute to lower cooling loads which represent a significant portion of a facility's energy consumption. White reflective surfaces in containment can improve lighting efficiency while sealed cabinets reduce airflow recirculation and improve cooling performance. Technologies such as aisle containment and controlled airflow pathways can help reduce energy use associated with cooling.

Designed to be built



How modern construction methods improve cost, programme and quality

BY ZAC POTTS, SUDLOWS, ASSOCIATE DIRECTOR (HEAD OF SUSTAINABLE ENGINEERING AND INNOVATION)

HEADING into 2026, data centres, the critical engine rooms of modern infrastructure, are now needed, in various forms, more than ever. The AI era adds real pressure and definite challenges, but there's also vast opportunity.

The data centre pipeline for 2026-8 is strong, but with the size and scale we need to deliver, in the target timeframe, it's certainly challenging. When a month's delay can equate to millions in lost opportunity, having certainty and predictability on programme is arguably as important as the technical performance.

At the recent Data Centre Alliance 10x10 Event at PowerEx London, I talked about how the right approach to design is one of the most powerful tools in risk mitigation and speed to market.

Listening to the other presentations and speaking with other attendees, it was clear; although size, region, and type of facility affect the specifics, the need for a shift in approach is across the board. The next wave of data centres fundamentally needs a new approach. The principle I spoke about – “designed to be built” – focussed on decisions

taken in design tangibly affecting the risk profile of the project, the speed of delivery, and the resilience in the design to accommodate change. By the end of Stage 4 design, the risk profile and programme is baked in, opportunities for proactivity reduce, and reactivity dominates.

Design needs to consider the overall project performance as M&E, civil, structural and architectural engineering are inextricably linked, both to one another and to the delivery programme, and the cost and risk profiles. Value engineering is a misnomer when it comes at the expense of programme. A design which relies on a single specific manufacturer or product is a high risk at procurement, whilst a backend heavy commissioning and testing period only identifies issues in the final hour.

Incorporating little to no flexibility, we risk becoming obsolete much earlier than intended. The solution to these issues is to design in a way which sees the risks and overall needs, and prioritises actual delivery.

Design not only dictates scope and cost, but influences risk profile, construction sequencing opportunities,

and interdependencies – driving the length and certainty of the critical path.

By embracing modularity and productisation, design and delivery start earlier, and become more controlled. In my experience, repeatable blocks bring not only efficiencies but a needed level of predictability.

Designing to be vendor agnostic and interoperable with various technology, may require slightly more space or lose out on the degree of optimisation which can only be delivered with fixed elements, but the benefits massively outweigh the costs.

Ultimately, the risk of insurmountable issues at the point of procurement is both unacceptable and entirely avoidable.

Integrating a delivery mindset into the design phase, as opposed to allowing one to flow blindly into the other, we don't just build better; we build faster and with more certainty. It's these foundations which I believe critically allow for a more scalable delivery model, ultimately converting the next few years of ambitious growth into reality.



The hidden energy asset inside every data centre



As demand for digital services continues to rise, data centres are facing growing scrutiny over their energy consumption. While much of the industry focus has been placed on IT efficiency and renewable energy procurement, cooling infrastructure remains one of the largest contributors to operational power demand particularly in high-density, mission-critical environments.

BY ANDY LEWIS, GENERAL MANAGER | DATA CENTRES AT WEATHERITE

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. As facilities scale to support AI workloads, cloud services, and edge computing, improving how thermal energy is managed is becoming increasingly important.

One area now gaining renewed attention is heat recovery not simply as a sustainability initiative, but as a practical strategy for improving overall energy efficiency and long-term operational resilience.

Reframing heat as a resource

Heat recovery involves capturing thermal energy that would otherwise be rejected and repurposing it for

productive use. In data centres, this heat is generated continuously by IT equipment and extracted by cooling systems to maintain stable operating conditions.

Traditionally, cooling strategies have been designed to remove heat as efficiently as possible and expel it from the site. However, as energy costs rise and carbon reduction targets tighten, this linear approach is being re-evaluated.

Because data centres operate 24/7, they produce a consistent and predictable stream of low-grade heat. This reliability makes them particularly well suited to recovery strategies, especially when compared to intermittent industrial processes or seasonal heat loads.

The scale of the opportunity

Recent research highlights the significance of this opportunity, particularly in urban environments. An AECOM-commissioned study examining London’s data centre estate found that the waste heat currently being expelled into the atmosphere could, if effectively captured and redistributed, theoretically provide enough thermal energy to heat hundreds of thousands of homes each year.

While not all of this heat is practically recoverable, the findings underline the scale of untapped potential within dense data centre clusters.

For operators, this shifts the conversation. Heat recovery is no longer just about incremental efficiency gains within a single facility; it is about



how data centres can play a more active role in wider energy systems.

Technical challenges and design considerations

Despite its benefits, heat recovery is not without complexity. Many existing data centres were designed around heat rejection rather than reuse, meaning retrofit projects must carefully consider temperature levels, system compatibility, space constraints, and cost.

Effective recovery depends on several technical factors:

- The temperature at which heat is available
- The method used to transfer and, where necessary, upgrade that heat
- Integration with building management systems
- Alignment between heat supply and demand profiles

Modern HVAC solutions are increasingly addressing these challenges. Indirect air cooling systems, hybrid architectures, and modular plant designs make it easier to incorporate heat recovery without compromising resilience or uptime. When paired with intelligent controls, recovery systems can operate dynamically, responding to IT load and external conditions to maximise efficiency.

Practical applications in and beyond the data centre

Heat recovery can be deployed in a range of ways depending on site layout and location.

Within the data centre campus, recovered heat can be used for space heating in offices, control rooms, and ancillary areas, as well as for domestic hot water pre-heating. In multi-building environments, heat can be shared with adjacent facilities that have compatible thermal demands. In urban locations, data centres are increasingly being viewed as anchor loads for district heating networks, exporting surplus heat to residential or commercial developments nearby. This approach not only improves overall energy utilisation but also strengthens relationships with local authorities and planning bodies.

Supporting compliance and ESG objectives

Heat recovery aligns closely with regulatory and sustainability



frameworks such as BREEAM, LEED, and the EU Code of Conduct for Data Centres. It also supports UK reporting requirements including SECR by delivering measurable reductions in Scope 1 and Scope 2 emissions. From an ESG perspective, heat recovery enables operators to demonstrate tangible operational decarbonisation rather than relying solely on offsetting strategies.

Forward planning for a heat-constrained future

As data centre capacity continues to expand, forward planning will be critical. Decisions made early in the design process around cooling architecture, plant layout, and system integration will determine whether facilities can adapt to future efficiency, regulatory, and sustainability requirements.

Heat recovery is most effective when it is considered at design stage, even if it is not immediately deployed. Designing

systems with future reuse in mind allows operators to preserve flexibility, minimise disruption, and unlock value as local heat demand, policy incentives, or infrastructure evolve.

This long-term approach requires collaboration with HVAC specialists who understand both mission-critical resilience and future thermal strategy. With more than 50 years of experience in HVAC innovation, Weatherite has seen the industry evolve from purely mechanical cooling solutions to highly integrated, digitally controlled systems supporting some of the most demanding environments in the UK.

That perspective is increasingly important. As expectations around energy performance continue to rise, data centres will be judged not only on uptime and efficiency today, but on how well they are prepared for the energy challenges of tomorrow. Heat recovery, when embedded within a forward-thinking design strategy, offers a practical way to turn a long-standing challenge into a lasting advantage.

Key Takeaways

- Data centres generate large volumes of consistent waste heat that can be repurposed to improve overall energy efficiency.
- Research from AECOM highlights the significant, largely untapped potential of heat recovery in urban data centre clusters.
- Modern HVAC technologies and intelligent controls are making heat recovery more achievable without compromising resilience.
- Forward planning and design-stage thinking will determine how effectively heat recovery becomes a standard feature of future-ready data centres.

Heat recovery involves capturing thermal energy that would otherwise be rejected and repurposing it for productive use. In data centres, this heat is generated continuously by IT equipment and extracted by cooling systems to maintain stable operating conditions