DCS
DATACENTRE SOLUTIONS
DEVELOPING DIGITAL INFRASTRUCTURE IN A HYBRID WORLD

ISSUE IV 2025

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

HTTPS://DATACENTRE.SOLUTIONS

DATA CENTRES AND REFRIGERANT USE: A CLIMATE CHALLENGE

Power protection

that's out of this world.

Galaxy VS UPS

Huge power in a small footprint...

The Galaxy™ VS from Schneider Electric is the best-in-class power protection that innovates against downtime. It is an efficient, modular, simple-to-deploy 3-phase UPS that delivers intelligent power protection to edge, small, and medium data centres, as well as other business-critical applications

- 99% efficient in ECOnversion™ mode
- Compact design, optimised footprint architecture
- Battery flexibility, including Lithium-ion batteries
- New patented hybrid technology

Discover the Galaxy VS UPS

www.se.com/uk

Life Is On | Schneider Electric

CONTENTS

ISSUE IV 2025

DCS DATACENTRE SOLUTIONS

14 Data centres and refrigerant use: A climate challenge

Data centres are vital in today's digital world. As the demand for digital services continues to grow, data centres are under increased pressure to minimise their environmental impact

16 AI data centres are hungry for power – here's how to manage their appetite

The root causes of data centres' increasing power demand, and proposes the solutions that will keep data centres online and efficient as they scale up

18 How data centres can become heat heroes

he convergence of digital infrastructure and energy systems represents a significant opportunity for data centre operators to reduce costs, generate new revenue, and meet sustainability targets

20 Immediate Power Solutions (IPS): a vital component in modern DCs

As the global economy relies more on data centers, the need for reliable UPS systems to ensure uninterrupted operations has become increasingly critical

22 How densification is impacting data centres: is hybrid cooling the answer?

As the push for net zero carbon emissions accelerates, data centres are under increasing scrutiny for their energy consumption and environmental impact

26 Why the data centres powering tomorrow's AI will be liquid-cooled

The power demands of generative AI has placed the energy demands of data centres under the spotlight

28 Should you be using power or water to cool your data centre?

There isn't a simple solution to power vs water-based cooling, understanding how the data centre fits into the wider ecosystem will ensure the correct choice

30 From silos to synergy

The kaleidoscope effect of data centre digital twins

32 The AI Era: redesigning data centres for scale and speed

AI's insatiable thirst for computing resources is transforming infrastructure, and the industry is grappling with how to meet the demands

34 Rethinking sustainability

Time to look at the bigger picture and avoid simplistic thinking?

38 Failure is not an option as AI downtime will be the most expensive ever

AI scaling challenges for MEP equipment supply and operation – and how to cope with them

40 Super-TCXO enables new architecture for greater AI workload efficiency in data centres

AI workload efficiency takes centre stage.

The data centre trade association

DCA News

46 DCA update
An Introduction from DCA CEO Steve Hone

46 Advancing liquid spray cooling
By Matthew Thompson, MD Europe, Airsys

48 Running the rule over refrigeration
By Lawrence Leask, MD of Kaizen Energy Consultancy

50 Winner takes all – meeting the demands of complex, end-date critical construction programs
By Christian Goldsmith, Senior Program Director, Technology, Arcadis

52 Why we can't stop building data centres – but need to do it differently
By Matt Edgley, Chief Operating Officer at Datum Datacentres

53 From load to leadership: How data centres are redefining decarbonisation for industry
By Charles Johnston, Sustainable Energy First

44 Advanced networks for artificial intelligence and machine learning computing
In today's rapidly evolving technological landscape, the demand for artificial intelligence (AI) and machine learning (ML) capabilities is driving significant transformations across industries.

NEWS

06 AI is set to drive surging electricity demand from data centres

07 Data centre energy storage market is projected to reach $4.3 bn by 2034

08 Breakfast club debates legacy data centre issues

09 Global data centre market set for 'unprecedented' growth

10 Alarming deficiencies in security readiness

11 AI ambitions, data doubts: 60% of business leaders feel unprepared

12 AI will transform core business models within the next two years

06

DATACENTRE SOLUTIONS

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

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

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

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

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

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

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

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

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

Angel BUSINESS COMMUNICATIONS

AI is set to drive surging electricity demand from data centres

Offers the potential to transform how the energy sector works.

MAJOR new IEA report brings groundbreaking data and analysis to one of the most pressing and least understood energy issues today, exploring AI's wide range of potential impacts

Artificial intelligence has the potential to transform the energy sector in the coming decade, driving a surge in electricity demand from data centres around the world while also unlocking significant opportunities to cut costs, enhance competitiveness and reduce emissions, according to a major new report from the IEA.

The IEA's special report Energy and AI, out today, offers the most comprehensive, data-driven global analysis to date on the growing connections between energy and AI. The report draws on new datasets and extensive consultation with policy makers, the tech sector, the energy industry and international experts.

It projects that electricity demand from data centres worldwide is set to more than double by 2030 to around 945 terawatt-hours (TWh), slightly more than the entire electricity consumption of Japan today. AI will be the most significant driver of this increase, with electricity demand from AI-optimised data centres projected to more than quadruple by 2030.

In the United States, power consumption by data centres is on course to account for almost half of the growth in electricity demand between now and 2030.

Driven by AI use, the US economy is set to consume more electricity in 2030 for processing data than for manufacturing all energy-intensive goods combined, including aluminium, steel, cement and chemicals. In advanced economies more broadly, data centres are

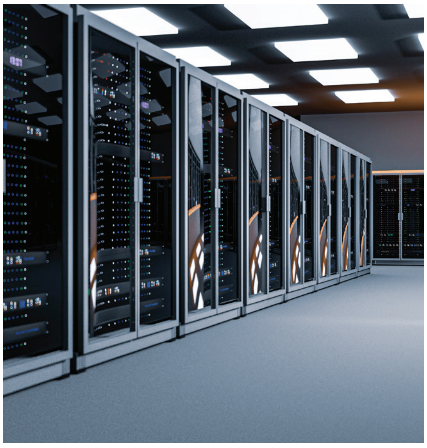

projected to drive more than 20% of the growth in electricity demand between now and 2030, putting the power sector in those economies back on a growth footing after years of stagnating or declining demand in many of them.

A diverse range of energy sources will be tapped to meet data centres' rising electricity needs, according to the report – though renewables and natural gas are set to take the lead due to their cost-competitiveness and availability in key markets.

"AI is one of the biggest stories in the energy world today – but until now, policy makers and markets lacked the tools to fully understand the wide-ranging impacts," said IEA Executive Director Fatih Birol. "Global electricity demand from data centres is set to more than double over the next five years, consuming as much electricity by 2030 as the whole of Japan does today. The effects will be particularly strong in some countries. For example, in the United States, data centres are on course to account for almost half of the growth in electricity demand; in Japan, more than half; and in Malaysia, as much as one-fifth."

The report emphasises the significant uncertainties that remain, from the macroeconomic outlook to how quickly AI will be adopted. It also notes questions over how capable and productive AI will become, how fast efficiency improvements will occur, and whether bottlenecks in the energy sector can be resolved.

AI could intensify some energy security strains while helping to address others, according to the report. Cyberattacks on energy utilities have tripled in the past four years and become more sophisticated because of AI.

At the same time, AI is becoming a critical tool for energy companies to defend against such attacks. Another energy security concern relates to the expanding demand for critical minerals used in the equipment in the data centres that power AI. The report provides first-of-its-kind estimates of demand from data centres for critical minerals, whose global supply is today highly concentrated.

While the increase in electricity demand for data centres is set to drive up emissions, this increase will be small in the context of the overall energy sector and could potentially be offset by emissions reductions enabled by AI if adoption of the technology is widespread, according to the report.

Additionally, as AI becomes increasingly integral to scientific discovery, the report finds that it could accelerate innovation in energy technologies such as batteries and solar PV.

According to the report, countries that want to benefit from the potential of AI need to quickly accelerate new investments in electricity generation and grids, improve the efficiency and flexibility of data centres, and strengthen the dialogue between policy makers, the tech sector and the energy industry.

Data centre energy storage market is projected to reach $4.3 bn by 2034

The Data Centre Energy Storage Market is set to grow from its current market value of more than $2 Billion to over $4.3 Billion by 2034; as reported in the latest study by Global Market Insights, Inc.

AS THE DEMAND for digital services surges, data centers are under increasing pressure to adopt efficient energy storage solutions that ensure uninterrupted power supply and optimize energy consumption. Businesses are moving away from traditional power grids, focusing instead on advanced storage technologies that align with sustainability goals and enhance operational resilience.

With growing concerns over carbon footprints and rising energy costs, the push for renewable energy integration is stronger than ever. Companies are leveraging cutting-edge storage systems to maintain uptime, reduce electricity expenses, and support the shift toward a greener digital infrastructure.

The rapid expansion of cloud computing, artificial intelligence, and big data analytics has intensified the need for scalable energy storage solutions. Digital transactions and data processing continue to escalate, increasing dependency on energy-efficient systems that can sustain high-performance computing. Organizations worldwide are investing in innovative energy storage technologies to improve power reliability and minimize the risk of downtime.

As regulatory frameworks emphasize sustainability, data centers are prioritizing eco-friendly energy solutions to remain competitive in the evolving market. The industry's transition toward hybrid energy storage, incorporating lithium-ion batteries, flywheels, and other energy-efficient systems, further underscores the shift toward a resilient, sustainable infrastructure. With the continued advancement of power management systems, data centers are optimizing energy efficiency while maintaining high-speed computing performance.

Market segmentation by data center size includes small, medium, and large facilities. In 2024, large data centers accounted for 46% of the market share, with revenue projections reaching USD 2 billion by 2034. High-capacity data centers depend on sophisticated energy storage solutions to ensure seamless operations, preventing disruptions in critical computing environments. As energy demands rise, businesses are integrating advanced storage technologies to enhance system resilience, reduce power failures, and optimize energy efficiency.

These large-scale centers require robust storage infrastructure to support increasing data loads, driving significant investments in energy storage systems. The data center energy storage market is further categorized by application, covering banking, energy, government, healthcare, manufacturing, IT, and colocation services. In 2024, colocation centers held a 30% market share, reflecting the growing preference for shared data storage facilities. Businesses utilizing colocation services prioritize continuous power availability, making energy storage solutions an essential component of their operational strategy. The increasing reliance on third-party data storage has intensified the demand for energy-efficient infrastructure, compelling colocation providers to implement cutting-edge power management solutions. Advanced energy storage technology ensures cost-effective operations while maintaining reliability, further driving adoption across the sector.

The US data center energy storage market accounted for 80% of the global share in 2024, solidifying its position as a major industry driver. The region's leadership in large-scale infrastructure projects has accelerated the adoption of advanced storage solutions as businesses seek more flexible and cost-effective energy consumption strategies.

Companies across North America are actively enhancing energy efficiency initiatives to meet sustainability targets, reducing carbon emissions, and improving power reliability. The strong focus on eco-friendly infrastructure and continuous investment in energy storage technologies have positioned the region at the forefront of market growth.

Breakfast club debates legacy data centre issues

The latest BCS Breakfast Club meeting was a 'full house' with over 30 participants which included leading industry investors, operators, designers and consultants.

SIMON HARRIS, Director of Critical Infrastructure at BCS, kicked off the meeting by inviting the room to share their views and experience around the subject of what to do with the considerable stock of legacy data centres to be found across the UK and Europe and the options available to ensure they remain relevant in this age of accelerating digital transformation. This was followed by a lively discussion with topics that included acquisition, refurbishment, upgrading, disposal and demolition of these vital parts of our digital infrastructure.

The Group discussed the benefits of refurbishing existing facilities which included sustainability, resource utilisation and speed to market. It was felt that the commonly experienced risks of site clearance and new construction has already been addressed and that in some cases the legacy facilities have underutilised power availability with consequential missed expansion and revenue opportunities. In addition, smaller former enterprise sites can be in good locations for edge deployments and finally there are savings to raw material resource utilisation and consequential carbon emissions when compared to new build, - an important ESG consideration which was discussed in further depth later in the session.

Several participants felt that refurbishment was a much more viable option for existing live operations than to demolish and rebuild because of the importance of assessing how existing customers would be best served alongside the need to honour existing contracts. As a result, it was felt that some organisations are looking to upgrade purely to support their clients and address concerns that aging equipment may cause outages rather than to upgrade to attract new ones and become AI ready. In short, they are wanting to squeeze out more efficiency rather than looking at future scalability. However, in many cases an upgrade and refresh to critical infrastructure could liberate trapped electrical capacity for deployment to serve higher density and growing IT loads, for example through UPS replacement or cooling solution changes. These types of interventions will be more easily accommodated in Tier III facilities having two concurrently maintainable power and cooling paths, although the work will require careful planning and right first time execution. Nevertheless, such a solution overcomes the power availability challenges and takes the facility down a path towards better PUE performance.

Amongst the group there was agreement that any form of major upgrade or refurbishment has risks as it is done in a live environment without impacting existing operations. There were several shared stories of complex, unplanned issues that were uncovered mid project due to a lack of information! However, it was noted that a number of legacy sites are located on prime land with good power connections so there is real value in refurbishing them if the risk can be managed although some felt strongly that it was simpler and more cost effective to build new facilities and that these were easier to manage.

Others were concerned about the lack of experienced people that are needed for a complicated, multifaceted refurbishment project. This was perhaps unsurprisingly given the well documented skills shortage and it was clear from all the participants that this is a problem throughout the supply chain. Some felt that this was even worse in mission critical projects which required a high level of engineering experience. This was described as a 'dearth of expertise' and there was further discussion about how to attract new young people into the sector whilst acknowledging that some progress is being made through investment in UTC, STEM, apprenticeships and T levels.

There was a discussion about the sustainability benefits of modernising legacy data compared to building a new one where a substantial amount of the construction work involves the use of energy dense concrete and steel to such an extent that refurbishing an existing facility can save in the order of 70% - 80% of the carbon cost of a new build. However, for legacy sites, energy efficiency is often a significant challenge. Older data centres were not designed with energy conservation as a high priority, leading to excessive power consumption and higher operational costs. This is not only financially burdensome but also environmentally unsustainable. With the growing emphasis on green computing and corporate responsibility, organisations are under pressure to upgrade their facilities to be more energy efficient. A potential carbon tax for the industry by the UK government was discussed along with a warning that "it is coming, and it will be substantial – so get ready'.

It was agreed that the opportunity to make the most of the embedded carbon in legacy data centres may encounter a challenge as it is hard to measure and there is currently no industry standard although it is likely that the Government will implement some. The dynamics of Scope 4 emissions were also discussed and how these might impact the sector overall.

The session concluded with an agreement of the challenges of balancing cost, risk and sustainability alongside the needs of today, tomorrow and the future. It was felt that each facility needed to be reviewed on a case by case basis as every one is different and there are challenges across finance, engineering, logistics, risk management and access to resources

Global data centre market set for 'unprecedented' growth

Knight Frank has published its global data centres report revealing a surge in market expansion, with a projected 46% increase in data centre capacity worldwide by 2027. This equates to an additional 20,828 megawatts (MW).

THIS RAPID GROWTH, which has the potential to expand 177% by 2030, is underpinned by a substantial capital expenditure of £229 billion over the forecast period, reflecting the intensifying demand for digital infrastructure to support AI, cloud computing, and enterprise digital transformation.

Following a 36% drop in data centre transaction volumes in 2023 due to global interest rate hikes, the market rebounded in 2024, surging 118% to £24.5 billion across single-asset purchases, portfolio acquisitions, redevelopment opportunities, and development site sales.

Globally, the average real estate transaction value in the data sector space was £59 million in 2024, up 15% on the average transaction price in 2023, and up 44% on the pre-COVID transactions value average in 2019. Since 2019, average transaction value has grown at a compound-annual-growth-rate (CAGR) of 7.5%.

Regional Growth Highlights:
North America remains the dominant global market, with 11,638 MW in new capacity, reflecting a 54% growth rate and a staggering £128 billion in capital being deployed to support this expected growth. The region benefits from a combination of homegrown hyperscale dominance, increasing enterprise colocation demand, and strategic expansion into emerging secondary markets.

Europe, Middle East & Africa (EMEA) is set to expand by 4,529 MW (44%), requiring a £49.8 billion investment. European markets are experiencing a shift towards secondary hubs such as Milan and Madrid, primarily driven by power constraints in core markets like Frankfurt and London.

Asia-Pacific (APAC) is forecast to see

a 4,174 MW (32%) increase, supported by a £45.9 billion investment. APAC remains a highly diverse market, with significant development in both established hubs like Tokyo and emerging locations such as Johor, Malaysia, where hyperscalers seek alternative expansion opportunities.

Key Markets Driving Expansion:
Ashburn, USA: The world's largest data centre hub will grow by 2,428 MW (58%), backed by £26.7 billion targeting this market. Despite power availability challenges, Northern Virginia remains the principal destination for hyperscalers and colocation providers.

Phoenix, USA: One of the fastest-growing markets, with a 126% surge (1,109 MW), attracting £12.2 billion in investment. The city's appeal is fuelled by its scalable land options, business-friendly environment, and strong connectivity infrastructure.

London, UK: Retaining its status as a leading European data centre market, London is set to expand by 480 MW (36%), with £5.3 billion of investment. However, ongoing power constraints in established submarkets is encouraging development in outer London and secondary UK cities.

Milan, Italy: The standout European market with a remarkable 168% growth rate (310 MW), requiring £3.4 billion in investment. Milan's rise is indicative of a broader shift in European data centre expansion towards new, less congested hubs.

Tokyo, Japan: A key APAC hub, poised for a 25% increase (295 MW) attracting £3.2 billion. Japan's strategic location, stable power grid, and increasing demand for cloud services continue to drive growth.

Johor, Malaysia: Emerging as a major data centre hotspot with an 85% growth rate (335 MW), underpinned by £3.7 billion in investment. Johor's proximity to Singapore, combined with attractive incentives, is establishing it as a viable alternative for hyperscale expansion.

Stephen Beard, Global Head of Data Centres at Knight Frank said, "The global data centre industry is undergoing rapid transformation, with hyperscaler and colocation providers prioritising markets that offer access to power, robust connectivity, and a favourable regulatory environment.

We're increasingly seeing sustainability considerations shaping investment strategies, with an increasing focus on renewable energy adoption and energy-efficient design. Real estate investors and developers are positioning themselves to capitalise on this demand, with an emphasis on acquiring strategically located land and securing long-term power agreements.

"As global capital races to capture the next wave of digital infrastructure growth, the competition for prime development sites, particularly in power-constrained locations, will intensify. Industry stakeholders must navigate regulatory complexities, power availability concerns, and sustainability requirements to remain competitive in this high-growth sector.

"Operators, investors, policymakers, and partners, each have a role to play in shaping this future. The task ahead is to build infrastructure that not only supports innovation but also safeguards sustainability, security, and equity."

Alarming deficiencies in security readiness

Cybersecurity readiness remains alarmingly low as only 4% of organizations worldwide have achieved a mature level of readiness.

ACCORDING to Cisco's 2025 Cybersecurity Readiness Index, only 4% of organizations worldwide have achieved the 'Mature' level of readiness required to effectively withstand today's cybersecurity threats. This is a slight increase from last year's Index, in which 3% of organizations worldwide were designated as Mature. This demonstrates that despite a slight improvement from last year, global cybersecurity preparedness remains low as hyperconnectivity and AI introduce new complexities for security practitioners.

AI is revolutionizing security and escalating threat levels, with nearly 9 in 10 organizations (86%) facing AI-related security incidents last year. However, only 49% of respondents are confident their employees fully understand AI related threats, and 48% believe their teams fully grasp how malicious actors are using AI to execute sophisticated attacks. This awareness gap leaves organizations critically exposed.

AI is compounding an already challenging threat landscape. In the last year, nearly half of organizations (49%) suffered cyberattacks, hindered by complex security frameworks with disparate point solutions. Looking forward, respondents view external threats like malicious actors and state-affiliated groups (58%) as more significant to their organizations than internal threats (42%), underscoring the urgent need for streamlined defense strategies to thwart external attacks.

"As AI transforms the enterprise, we are dealing with an entirely new class of risks at unprecedented scale - putting even more pressure on our infrastructure and those who defend it," said Cisco Chief Product Officer Jeetu Patel. "This year's report continues to reveal alarming gaps in security readiness and a lack of urgency to address them. Organizations must rethink their strategies now or risk becoming irrelevant in the AI era."
2025 Cisco Cybersecurity Readiness Index: Cybersecurity Readiness Remains Flat as AI Transforms the Industry

The Index evaluates companies' readiness across five pillars—Identity Intelligence, Network Resilience, Machine Trustworthiness, Cloud Reinforcement, and AI Fortification—and encompassing 31 solutions and capabilities. Based on a double-blind survey of 8,000 private sector security and business leaders in 30 global markets, respondents detailed their deployment stages for each solution. Companies were then categorized into four readiness stages: Beginner, Formative, Progressive, and Mature.
Findings

The lack of cybersecurity readiness globally is alarming as 71% of respondents anticipate business disruptions from cyber incidents within the next 12 to 24 months. Further:
AI's Expanding Role in Cybersecurity: An impressive 89% of organizations use AI to understand threats better, 85% for threat detection, and 70% for response and recovery, underscoring AI's vital role in strengthening cybersecurity strategies.

GenAI Deployment Risks: GenAI tools are widely adopted, with 51% of employees using approved third-party tools. However, 22% have unrestricted access to public GenAI, and 60% of IT teams are unaware of employee interactions with GenAI, underscoring major oversight challenges.

Shadow AI Concerns: Sixty percent of organizations lack confidence in detecting unregulated AI deployments, or shadow AI, posing significant cybersecurity and data privacy risks.

Unmanaged Device Vulnerability: Within hybrid work models, 84% of organizations face increased security risks as employees access networks from unmanaged devices, further exacerbated by using unapproved Gen AI tools.

Investment Priorities Shift: While 96% of organizations plan to upgrade their IT infrastructure, only 45% allocate more than 10% of their IT budget to cybersecurity (down 8% year-over-year), emphasizing a critical need for more focused investment in comprehensive defense strategies, which is incredibly important as threats are not slowing.

Complex Security Postures: Over 77% of organizations report that their complex security infrastructures, dominated by the deployment of more than ten point security solutions, are impeding their ability to respond swiftly and effectively to threats.

Talent Shortage Impedes Progress: A staggering 86% of respondents identify the shortage of skilled cybersecurity professionals as a major challenge, with more than half reporting more than ten positions to fill.

To tackle today's cybersecurity challenges, organizations must invest in AI-driven solutions, simplify security infrastructures, and enhance AI threat awareness. Prioritizing AI for threat detection, response, and recovery is essential, as is addressing talent shortages and managing risks from unmanaged devices and shadow AI.

AI ambitions, data doubts: 60% of business leaders feel unprepared

60 per cent of business leaders are unsure about their organisation's readiness for data and AI, signalling a critical obstacle to progress at a time when AI adoption is accelerating across industries, according to research.

THE REPORT, conducted by the Business Performance Innovation Network, Growth Officer Council and EncompaaS, surveyed 170 global business decision makers to observe business model variances around AI and data.

While optimism around Generative AI remains high, the research uncovered a stark disparity between expectation and readiness. Despite 79 per cent of leaders anticipating a competitive advantage from GenAI in the next 18 months, just 13 per cent described themselves as "extremely confident" in their organisation's data-AI maturity. Key barriers include data accuracy and reliability (69%), AI integration and implementation challenges (68%), and ethical concerns including governance and trust (58%).

The findings point to a growing disconnect between businesses' desire to be data-driven and their current ability to manage and leverage data effectively—an issue that could stall innovation if not urgently addressed. As companies look to unlock the potential of AI for everything from forecasting to customer experience, the need for robust data infrastructure has never been more apparent.

Sachin Agrawal, Managing Director for Zoho UK, commented: "Artificial Intelligence is already having a transformative impact on business processes such as data analysis, forecasting and customer experience. However, maximising these benefits requires a foundation of high-quality, well-governed data.

To truly innovate, businesses must ensure that the tools they use are intuitive, integrated, and accessible to all employees—not just technical teams. Simplifying data management and improving user experience is key to enabling AI to flourish across every department."

Our own Zoho Digital Health Study shows that 46 per cent of businesses with good digital health are already seeing AI pay dividends. The message is clear: organisations that invest in data governance, quality, and workforce upskilling will be best placed to lead responsibly and competitively in the AI era."

Machine Learning projected to grow 40% faster than AI industry average by 2031

Despite economic headwinds, tighter AI regulations, Trump-era tariff policies, and escalating global trade tensions, the AI industry continues to outperform nearly every other corner of the tech sector. By 2031, the market is on track to surpass a staggering $1 trillion in value, and 56% of that figure will come from machine learning, AI`s largest and fastest-growing segment.

According to data presented by AltIndex.com, machine learning is projected to grow 40% faster than the AI industry average by 2031.

Machine Learning Market Set to Surge 440% by 2031

While the entire AI industry has experienced explosive growth in recent years, the rise of machine learning has been nothing short of extraordinary, and that momentum shows no signs of slowing. This technology has transformed industries from finance to healthcare and continues to attract record investments, showing investors believe in its long-term potential.
In Q1 2025 alone, companies and startups working in the machine learning space raised a mind-blowing $54.8 billion, the highest quarterly figure in market history. With the surge

of VC investments, machine learning continues its meteoric growth, far outpacing other AI sectors and the broader market.

According to the latest Statista Market Forecast, the machine learning industry is expected to reach $105.4 billion in 2025, a $30 billion increase in just one year. But that's only the beginning. The market is projected to skyrocket by nearly 440%, hitting an impressive $568 billion by 2031. This triple-digit growth is even more impressive when compared to other AI sectors.

Statistics show machine learning will grow 40% faster than the industry average of 331% by 2031. Natural language processing will trail even more with a 277% increase, or 58% lower than machine learning's growth. AI robotics follows with a 316% jump, reflecting a 38% gap. Meanwhile, computer vision and autonomous sensor technologies lag far behind. Statistics show computer vision will grow 200% slower than machine learning, while autonomous and sensor technology is expected to fall behind by a staggering 300%, highlighting the huge difference in growth rates across AI's main sectors.

United States and China in a Close Race for Machine Learning Dominance
In global comparison, the United States represents a machine-learning superpower, but China is quickly closing the gap.

AI will transform core business models within the next two years

Endava has launched its latest research report "AI and the Digital Shift: Reinventing the Business Landscape".

THE REPORT delves into the impact of the digital shift and explores how organisations are navigating this change and preparing for the future.In partnership with Censuswide, Endava surveyed 350 business and technology leaders across the US, EMEA and APAC to uncover their views on the digital shift and the role AI plays in reshaping business models.

Key findings include:
The impact of the digital shift: 98% of respondents have seen their core business model impacted by the digital shift. 37% are exploring new markets, while 24% say their organisation's mission or purpose has evolved as a result.

AI's role in reshaping core business models: 97[1]% of business leaders surveyed believe AI will transform core business models within the next two years. 31% believe AI will drive their organisation into new markets, 31% say it will help them focus on a unique selling proposition (USP) to differentiate from competitors, and another 24% say AI hasl impacted their core purpose.
The pressure to deliver: Business and technology leaders are under pressure to demonstrate AI's value. More than half (53[2]%) express concerns that if AI projects fail, their stakeholders or board will lose confidence in AI's potential. Additionally, 18% cite a lack of clear ROI as one of the top challenges in implementing and scaling AI solutions.
AI for creativity is untapped: 45% of respondents say their company uses AI to automate processes and repetitive tasks, but only 23% are leveraging AI to design new products. This highlights the vast untapped potential of AI, suggesting that while it is currently utilised for routine tasks, it holds promise for driving creativity in the future.

The talent gap remains a major roadblock: 41% of organisations cite a lack of skilled AI talent as the biggest challenge in implementing and scaling AI solutions. Respondents also noted it as the second highest barrier to digital modernisation efforts (36%).

Technology leaders are setting the pace: Business and technology leaders surveyed have different perspectives on AI readiness, with technology leaders (96%2) feeling more confident than their business counterparts (83%2). 86[3]% of technology leaders say their organisation integrates AI within its workforce to drive innovation and productivity effectively, compared to 74%3 of business leaders.

AI will soon be indispensable: There's a clear consensus that AI is becoming essential to business success. On average, respondents believe AI will become indispensable to their organisation's survival by 2030[4]. More than 3 in 5 (62%[5]) say this will happen within the next five years, and 11% already consider AI as critical.

Endava CTO, Matt Cloke, commented: "Just as digital transformation shapes entire industries, AI is set to redefine how businesses operate, innovate and succeed. This isn't about surface-level tool enhancements; AI must be embedded at the core of an organisation to drive meaningful change."

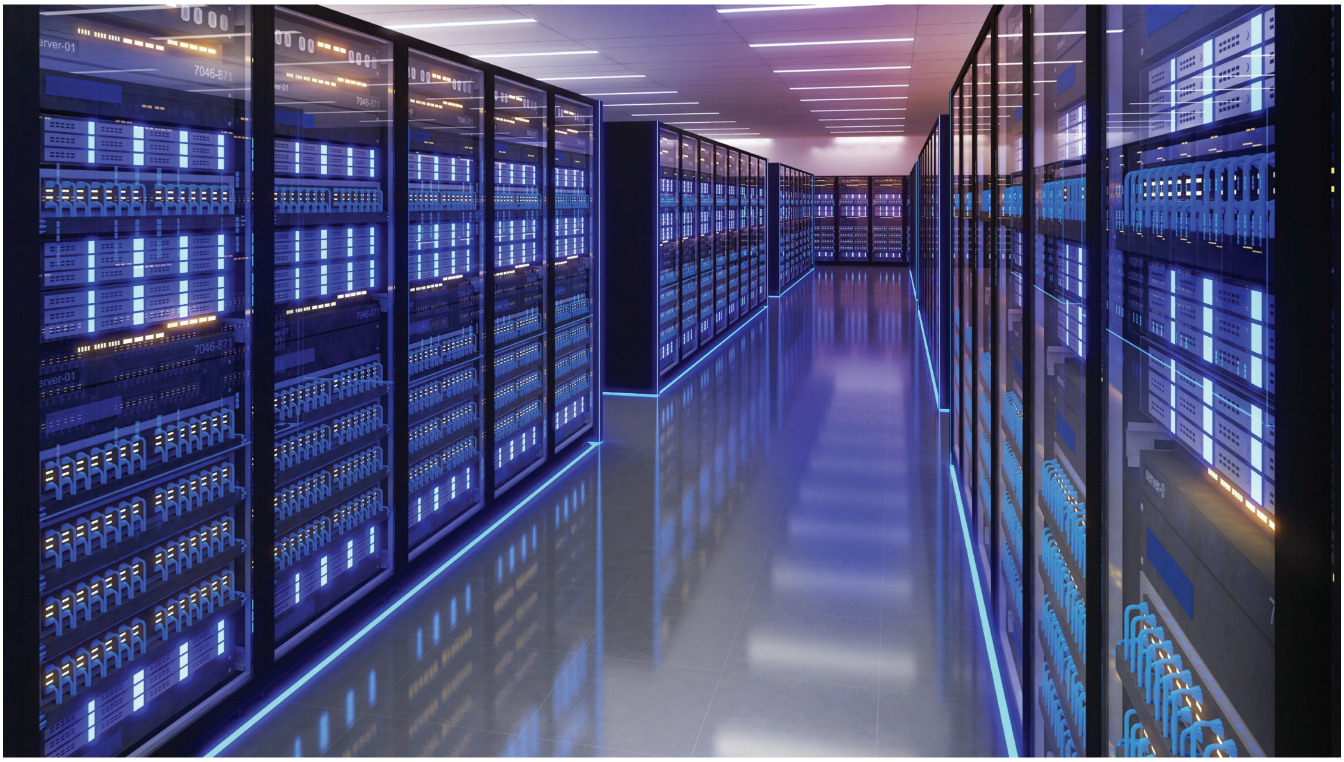

Data Centres and Refrigerant Use: A Climate Challenge

Data centres are vital in today's digital world. As the demand for digital services continues to grow, data centres are under increased pressure to minimise their environmental impact

BY A-GAS

THE DIGITAL infrastructure sat behind social media platforms, streaming services and cloud computing relies heavily on Heating, Ventilation, Air Conditioning, and Refrigeration (HVAC-R) systems to maintain optimal operating temperatures. But behind this infrastructure there is a complex challenge; these systems rely heavily on refrigerants that can pose a threat to the environment if they are not properly managed.

To help counter this, companies like A-Gas, through its dedication to recovery, reclamation and reuse of refrigerants are enabling the circularity of refrigerants across the data centre industry.

Enabling circularity

As the world faces increasing pressure to decarbonise, the transition to a circular economy, a model that keeps products and materials in use for as long as possible, has become an essential strategy for achieving environmental, economic, and operational sustainability.

This shift is especially urgent in high impact sectors such as Information Technology (IT) and HVAC-R.

Recovery and reclamation of refrigerants are the circular economy in action at A-Gas as we encourage the industry to make the most of the gases we already have. Collecting waste and re-using it is both economically and environmentally efficient. It overcomes the traditional linear model of make, use, and dispose which has dominated the way of business for the past 50 years. This is a particularly wasteful model and one that assumes there will be a never-ending supply of raw materials. As an industry faced with environmental pressures, the move towards a re-use approach is imperative.

By increasing circularity, we are making sure that these gases are not vented or leaked into the atmosphere, and through their re-use, we are also avoiding the need to produce the equivalent quantity of virgin refrigerants; this effectively doubles the emissions savings on offer as it reduces the emissions associated with new product manufacturing and shipping.

A-Gas and Lifecycle Refrigerant Management (LRM): A circular approach to refrigerants

Lifecycle Refrigerant Management (LRM) plays a vital role in helping mitigate climate change and its principle is a simple one:
"No kilogram or pound of refrigerant, once produced, should find its way into the atmosphere*.

Refrigerant recovery is a fundamental part of LRM; if refrigerants are not recovered, they cannot be reclaimed for further reuse. It is clear that recovering existing refrigerants for future re-use, or destruction, has a far greater impact on our climate than the substitution of lower global warming potential refrigerants alone.

To reduce environmental impact and improve efficiency, we must adopt refrigerant management practices that support circularity, including:

- **Refrigerant Recovery and Reuse:** One of the most impactful actions is to recover, reclaim, and reuse refrigerant gases rather than producing virgin refrigerants from raw materials.

- **Adoption of lower global warming potential refrigerants:** Switching to lower GWP refrigerants ensures compliance with the F-Gas Regulations and enables more efficient cooling technologies, helping reduce overall energy consumption.

Looking ahead

As sustainability goals grow more ambitious and regulations tighten, transitioning to lower global warming potential refrigerants while simultaneously strengthening refrigerant management must become a core part of decarbonisation strategies.

As an industry, we must transition to lower GWP refrigerants and ensure that potentially harmful legacy refrigerants are carefully managed to prevent their release into the atmosphere.

The recovery and subsequent reclamation or safe destruction of used refrigerants is key to mitigating the environmental impact of products that remain critical to providing digital services.

*LRM, the 90 Billion Ton Opportunity

AI data centres are hungry for power — here's how to manage their appetite

The root causes of data centres' increasing power demand, and proposes the solutions that will keep data centres online and efficient as they scale up.

BY ABB'S DATA CENTRES PORTFOLIO MANAGER, DANEL TURK

THE INTERNATIONAL ENERGY AGENCY (IEA) predicted in 2024 that global data centre electricity consumption would double to reach 1,000 terawatt hours (TWh) by 2026, equivalent to Japan's annual electricity consumption. The world's biggest tech companies plan to invest $1 trillion in new data centres over the next five years, primarily because of the sustained growth of AI and cloud services.

However, AI and machine learning models require far more computing power than traditional cloud computing functions, putting enormous pressure on local utility grids. For example, processing the average ChatGPT query demands 15 times more energy than a traditional web search.

Data centres are also expensive to operate, but downtime is even more costly. Every minute a data centre is offline can result in losses of tens of thousands of pounds. Consequently, maintaining efficiency and round-the-clock reliability is crucial to remain competitive, profitable and scalable.

So, data centre professionals face several challenges when looking to meet these demands cost-effectively and sustainably. Here are a few strategies that will enable you to overcome them.

Take advantage of AI

While AI is driving the increase in power demand, it might also help solve it via optimised cooling, predictive condition-based maintenance, data handling, and effective demand balancing. Connectivity plays a crucial role here. An AI-driven automation system can monitor the vast array of data points present in a medium to large-scale data centre. This capability provides operators with a comprehensive view of the data centre's performance, energy consumption, and asset health at any moment.

Operators can then use these insights to make targeted efficiency improvements. Consider, for instance, the cooling system. Typically, the upstream chiller and distribution system are viewed as separate entities, which means that efforts to improve the efficiency of one component might inadvertently reduce the efficiency of the other.

However, an automation system allows operators to understand their interdependencies. This holistic perspective empowers them to make informed decisions to improve the entire system's efficiency.

Plan modular

When constructing or expanding a data centre, consider adopting a modular approach. Modularity allows for sustainable scaling, meeting power and availability demands while simplifying specification and installation processes. Instead of building the entire facility simultaneously, you can develop it in stages. For example, a planned 200-megawatt (MW) data centre can be built in 20MW sections, onboarding customers as demand grows. This method is more cost-effective, avoiding excess capacity.

It also minimises downtime and risk, especially when using prefabricated solutions like skids and eHouses, which are assembled off-site and factory-tested before delivery. Purchasing a single prefabricated product, rather than multiple components requiring assembly, saves time and money. Local utilities and governments often prefer this approach, as the gradual growth of a data centre is more likely to gain approval than a full-scale new construction.

Extend modularity to your UPS

A modular approach is also highly effective in a UPS setup. Many data centres still use traditional monolithic UPS systems, where all components are built into single blocks. This design limits capacity adjustments without replacing the entire system or adding another UPS. In contrast, modular UPS systems comprise multiple smaller, interchangeable, autonomous modules working together in one frame. For instance, a one-megawatt (MW) modular UPS can be constructed from four parallel 250-kilowatt (kW) modules.

While monolithic UPS systems often have lower initial costs, modular setups offer better TCO and availability due to their inherent redundancy. Operators can specify a single module rather than an entire monolithic UPS block. To illustrate, achieving N+2 redundancy for a 1 MW data centre with a 1.5 MW modular UPS frame can save 30% compared to N+N redundancy using two 1 MW monolithic blocks.

> An AI-driven automation system can monitor the vast array of data points present in a medium to large-scale data centre. This capability provides operators with a comprehensive view of the data centre's performance

Modular systems are often lighter, facilitating installation on raised floors or rooftops. They also reduce the need for additional power equipment like cables and switchgear. In the scenario above, two parallel monolithic UPS would require two input and two output switches, while a modular setup needs only one of each.

Moreover, modular UPS systems are more energy efficient. During periods of low demand, modern modular setups can automatically switch modules to sleep or economy mode, reducing energy consumption.

Consider shifting to medium voltage

Data centres are becoming bigger and will continue to do so. This rise in power demand is making medium-voltage (MV) equipment increasingly more viable and cost-effective.

An MV Uninterruptible Power Supply (UPS) can deliver power protection to the entire data centre, not just to the racks like a low-voltage (LV) UPS would. They also bring notable energy-efficiency gains, resulting in long-term cost savings and a more attractive total cost of ownership (TCO).

MV UPSs also have a modular installation capability. This avoids relying on numerous LV UPSs that require regular service and maintenance, making MV UPSs more cost-effective.

For instance, you can parallel wire ten 2.5 MW UPS blocks to form a collective 25 MW system. Such configurations bolster the rapid deployment of additional units, increasing overall system capacity without introducing added complexity. Additionally, the lower currents at MV levels require cables with smaller cross-sections, reducing initial costs and improving TCO.

A positive contribution

Power availability will always be a pain point for data centre operators. By considering these strategic approaches, data centre managers can meet increasing power demands while upholding sustainability commitments and managing costs. As AI and digital services expand, these practices will be crucial in ensuring data centres positively impact the global energy landscape.

How data centres can become heat heroes

The convergence of digital infrastructure and energy systems represents a significant opportunity for data centre operators to reduce costs, generate new revenue, and meet sustainability targets.

BY DANIEL BALL, SENIOR BUSINESS DEVELOPMENT MANAGER AT SSE ENERGY SOLUTIONS

IN TODAY'S digital economy, data centres serve as the backbone of our interconnected world, processing and storing vast amounts of information that power everything from streaming services to financial transactions.

However, these digital powerhouses come with significant energy demands and environmental challenges. As the sector continues its rapid expansion, innovative solutions are emerging to address both the operational costs and carbon footprint of these facilities. One of the most promising approaches is the integration of data centres into heat networks, creating a symbiotic relationship that delivers substantial economic and environmental benefits.

The heat challenge in modern data centres

Data centres generate enormous amounts of heat as a by-product of their operations. The computing equipment packed into these facilities requires constant cooling to prevent overheating and ensure optimal performance.

Traditional cooling methods are energy-intensive, with cooling systems typically accounting for up to 40% of a data centre's total energy consumption. This not only represents a significant operational cost but also contributes substantially to the facility's carbon footprint.

As data processing demands grow exponentially, so does the heat management challenge. The global data centre cooling market is projected to reach £20 billion by 2026, reflecting the scale of this issue.

Moreover, with the rise of high-density computing and artificial intelligence applications, which generate even more

heat per square metre, the need for efficient heat extraction solutions has never been more pressing.

Heat networks: Turning waste into resource

Heat networks, also known as district heating systems, distribute heat from a centralised source to multiple buildings or facilities through insulated pipes. Traditionally, these networks have relied on dedicated heat generation plants, but data centres present an opportunity to repurpose what would otherwise be wasted energy.

By capturing excess heat from data centre operations and redirecting it into heat networks, facility operators can transform a costly waste product into a valuable resource. This recovered heat can be used to warm nearby buildings, supply hot water, or support industrial processes, creating new revenue streams while reducing the environmental impact of both the data centre and the broader community it serves.

Economic benefits for data centre owners

The financial case for connecting data centres to heat networks is compelling. Studies indicate that data centre operators can reduce their cooling costs by up to 30% through heat recovery systems. This represents significant savings, particularly for hyperscale facilities where cooling expenses can run into millions of pounds annually.

Beyond direct cost savings, data centres can generate new revenue by selling recovered heat to network operators or end users. Stockholm Data Parks estimates data centers can generate €0.5-1 million (equivalent of £420,000 to £840,000) in annual revenue per 10 MW of IT load by selling excess heat to district heating networks. These income streams have the added advantage of being relatively stable and predictable, diversifying revenue beyond traditional data services.

Additionally, heat network participation can extend the lifespan of cooling infrastructure by reducing the operational burden on these systems. With cooling equipment representing approximately 15% of initial capital expenditure for a typical data centre, any extension of this infrastructure's useful life translates to substantial financial benefits.

Carbon footprint reduction and sustainability gains

The environmental case for heat recovery is equally strong. By repurposing waste heat, data centres can dramatically reduce their carbon footprint—up to 97% reduction in certain applications, comparable to the emissions reduction achieved when switching from coal-fired power generation to wind power, as noted in recent renewable energy research.

> Despite the clear benefits, integrating data centres into heat networks presents several technical and logistical challenges

This carbon reduction potential is particularly valuable as data centre operators face increasing regulatory pressure and stakeholder expectations regarding sustainability. Major tech companies have committed to carbon-neutral or carbon-negative operations, and heat recovery represents a practical pathway to meeting these ambitious targets.

Furthermore, heat network participation helps address the growing scrutiny of data centres' resource consumption. As communities become more conscious of the environmental impact of these facilities, demonstrating positive contributions to local energy systems can strengthen a data centre's social license to operate and improve relationships with local authorities and residents.

Implementation challenges and solutions

Despite the clear benefits, integrating data centres into heat networks presents several technical and logistical challenges. These include matching heat supply with demand patterns, ensuring appropriate water temperatures for network distribution, and coordinating infrastructure development between different stakeholders.

Forward-thinking data centre operators are addressing these challenges through advanced heat recovery technologies and innovative partnerships. For example, high-efficiency heat pumps can upgrade low-temperature waste heat to the higher temperatures required by most heating systems. Meanwhile, collaborations between data centre developers, energy companies, and local authorities are creating the frameworks needed for successful integration.

Location strategies are also evolving, with new facilities increasingly being sited with heat network potential in mind. Urban and suburban locations close to areas with significant heat demand offer the greatest potential for efficient heat reuse, challenging the traditional preference for isolated data centre campuses.

The future of integrated energy systems

As we move toward more intelligent and interconnected energy systems, data centres are positioning themselves not just as consumers of resources but as active contributors to sustainable energy landscapes. Beyond heat networks, some facilities are exploring complementary approaches such as renewable energy generation, energy storage, and demand response capabilities.

The convergence of digital infrastructure and energy systems represents a significant opportunity for data centre operators to reduce costs, generate new revenue, and meet sustainability targets. For those willing to embrace this approach, the rewards extend beyond immediate financial gains to include enhanced reputation, regulatory compliance, and long-term business resilience.

In a world increasingly concerned with both digital connectivity and environmental sustainability, the integration of data centres into heat networks offers a compelling solution that addresses both priorities. For data centre owners and operators, the question is not whether to explore this opportunity, but how quickly they can implement the necessary technologies and partnerships to realise its benefits.

Immediate Power Solutions (IPS): a vital component in modern data centres

As the global economy relies more on data centers, the need for reliable UPS systems to ensure uninterrupted operations has become increasingly critical.

BY TOD HIGINBOTHAM, COO, ZINCFIVE

THE RAPID evolution of technology—spanning online services, consumer electronics, IoT, and AI—has amplified our reliance on digital infrastructure and emphasized the importance of uninterrupted data center operations. This technological surge is driving advancements in electrical infrastructure, particularly in energy storage, due to evolving needs. Data center operators face the challenge of addressing these increasing demands while also managing workplace safety, rising property expenses, and sustainability issues. Consequently, there is a noticeable shift towards backup power solutions that prioritize greater reliability, space efficiency, and environmental stewardship.

The growing pressure on data centers has brought Immediate Power Solutions (IPS) to the forefront as a crucial category of energy storage. IPS addresses the evolving digital infrastructure landscape by focusing on the immediate, high-rate power necessary for critical operations, distinguishing it from traditional Energy Storage Systems (ESS). Unlike ESS, which focuses on long-duration capacity, IPS is designed to deliver high-rate power instantly for short durations, emphasizing reliability and efficient space use.

When assessing mission-critical backup applications through the IPS framework, the suitability of different systems becomes clearer. For example, lead-acid batteries, which utilize one of the oldest battery technologies, are frequently viewed by many data center operators as a traditional and reliable choice for data center uninterruptible power supplies (UPS). However, given today's data center requirements for safety, reliability, sustainability, and space efficiency, more advanced battery technologies have now surpassed lead-acid in their suitability for the sector. One alternative considered by data center operators for their UPS systems is lithium-ion, an ESS battery. These batteries' energy density—characterized by their ability to release moderate amounts of energy over extended periods—makes them well-suited for applications such as electric vehicles and consumer electronics. However, data centers need a different capability: the ability to deliver rapid, high-power bursts to keep operations running during power outages until backup generators activate. While lithium-ion batteries are effective for sustained energy output, they may not meet the immediate high-power demands required for data centers. In contrast, IPS solutions are specifically designed to provide the rapid, high-power energy needed during these critical moments and do not have the same space and safety challenges associated with lithium-ion batteries.

Recent advancements in battery technology have introduced solutions specifically tailored for IPS needs. Notably, nickel-zinc batteries represent a breakthrough innovation in this field. Their enhanced power density allows them to deliver substantial bursts of energy rapidly while occupying less than half the space of conventional lead-acid systems. Their compactness and efficiency are particularly advantageous for data centers, as they align with the core objectives of IPS by optimizing space usage, supporting uninterrupted data center operations, and improving overall operational effectiveness and reliability.

IPS batteries can also contribute to UPS systems' reliability and safety. Unlike other batteries that may experience cell failures, batteries such as nickel-zinc maintain conductivity even if some cells are depleted, ensuring continuous operation. They also avoid thermal runaway and tolerate higher temperatures better than other types, adding an extra layer of safety and stability crucial for maintaining uninterrupted data center operations. As environmental concerns increasingly influence technological development, IPS is advancing to meet new sustainability standards as well. Nickel-zinc batteries, for example, are designed with sustainability in mind, offering lower lifecycle emissions and reduced resource consumption. They require fewer resources during production and utilize abundant, less environmentally taxing materials.

By adopting IPS for UPS, data centers not only achieve their sustainability goals but also benefit from high-performance energy storage solutions. This alignment with environmental objectives, combined with the efficient power delivery and space optimization inherent to IPS technologies, underscores the critical role of innovative solutions in advancing both operational effectiveness and ecological responsibility in modern data centers. Finally, transitioning to advanced IPS technologies is becoming increasingly seamless thanks to recent innovations in battery technology. These advancements have simplified the integration of new solutions into existing systems, allowing data centers to upgrade to more efficient and sustainable IPS options with minimal disruption. This streamlined process not only enhances operational efficiency but also ensures that data centers can quickly adapt to the evolving demands of modern infrastructure, positioning them for long-term success.

As the global economy relies more on data centers, the need for reliable UPS systems to ensure uninterrupted operations has become increasingly critical. This growing demand for dependable, space-efficient, and eco-friendly backup power solutions is driving the shift towards IPS. Their alignment with the challenges they're solving, as well as the seamless integration into existing systems, facilitates a smoother transition to advanced battery solutions, ensuring that data centers remain resilient and future-ready.

How densification is impacting data centres: is hybrid cooling the answer?

As the push for net zero carbon emissions accelerates, data centres are under increasing scrutiny for their energy consumption and environmental impact. Cooling systems, which account for a significant share of energy use, are a critical focus for sustainability efforts – and high-temperature fluid coolers, a key component of advanced chilled water systems, have emerged as transformative technology in this space.

BY GEORGE HANNAH, SENIOR GLOBAL DIRECTOR FOR CHILLED WATER SYSTEMS AT VERTIV

BY ENABLING heat recovery, reducing energy demand and supporting compliance with stringent environmental regulations, high-temperature fluid coolers are reshaping how data centres operate. Their integration into sophisticated cooling strategies positions them as a cornerstone of efficient infrastructure and a driver of the industry's path to net zero.

Enabling net zero strategies

The transition to net zero in the data centre industry is driven by

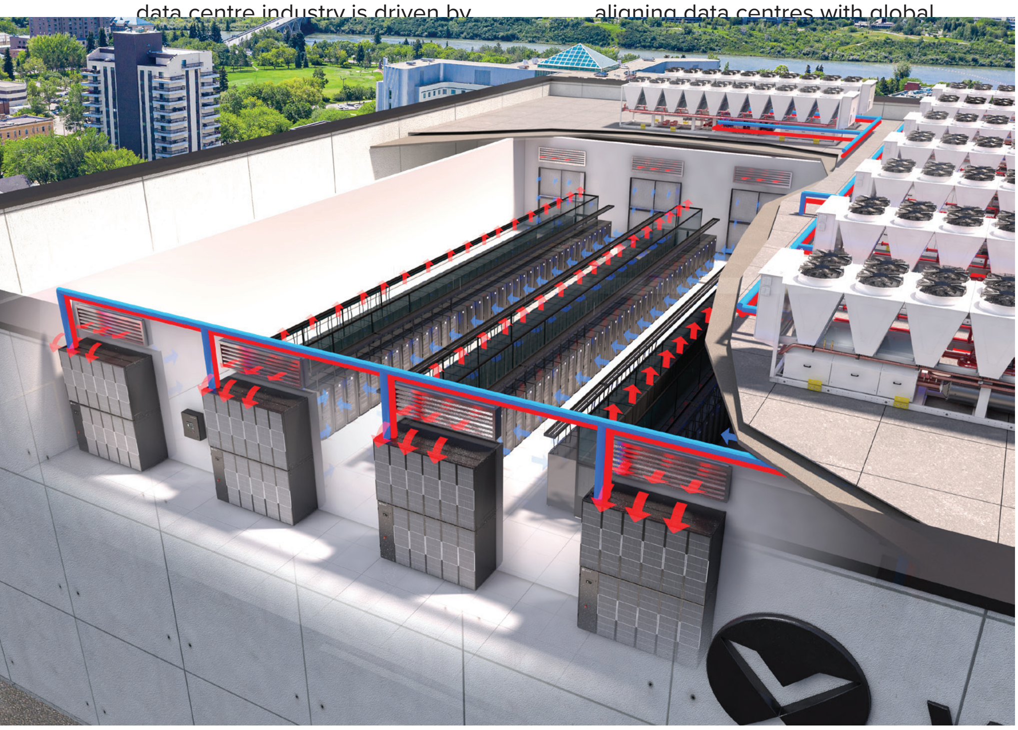

energy. Cooling systems, which have traditionally been seen as an operational burden, are now viewed as an opportunity for real innovation and impact. High-temperature fluid coolers are designed specifically to support these goals.

By operating efficiently at elevated water temperatures - up to 40°C supply and 50°C return, these systems reduce the energy required for cooling. This not only lowers operational costs but also decreases carbon emissions, aligning data centres with global sustainability standards such as the Science Based Targets initiative (SBTi). The use of heat recovery systems also allows operators to take their sustainability efforts further. Captured waste heat can be repurposed for district heating, reducing reliance on grid-sourced power, often linked to fossil fuel use.

The capability to reuse energy enhances the broader environmental impact of data centres, transforming them from isolated facilities into active contributors to decarbonisation.

The impact of regulations

Regulations around carbon emissions, energy efficiency, and refrigerant use are tightening globally. Frameworks like F-gas Regulation and local emissions reduction mandates are putting cooling systems under the microscope.

As data centres grow in scale and prominence, failing to meet these evolving standards risks reputational and operational setbacks.

High-temperature fluid coolers help facilities stay ahead of these regulatory demands. Many models now use very low Global Warming Potential (GWP) refrigerants, enabling compliance with stringent CO2e emissions standards. Additionally, their high efficiency and heat recovery capabilities position them as ideal solutions for facilities navigating stricter energy efficiency requirements.

For operators looking to design their facilities to be future-ready, these systems offer a clear path to compliance without sacrificing performance. The ability to meet current and anticipated regulations makes them an indispensable part of a forward-looking cooling strategy.

Unlocking the potential of heat recovery

Heat recovery is a growing priority for data centres aiming to reduce their environmental footprint. High-temperature fluid coolers excel in this area, removing the heat generated and repurposing it for secondary applications.

For example, waste heat can be used to support district heating networks, providing an energy-efficient alternative to traditional heating methods. It can also be used to thermally power local industrial processes, contributing to circular economy goals by maximising resource efficiency.

By integrating heat recovery into their operations, data centres can move beyond mere compliance and take on a leadership role in sustainability.

To demonstrate how effective their heat reuse is, operators should align with evolving sustainability metrics such as Energy Reuse Effectiveness (ERE) and Heat Recovery Efficiency (HRE). ERE considers how well energy is recovered most likely in the form of heat.

HRE measures how efficiently the recovered energy is reused. A higher HRE indicates greater efficiency in utilising waste heat. Metrics like these, that account for energy recovery will become increasingly important and chilled water solutions are central to achieving these goals.

High-temperature fluid coolers enable operators to achieve strong performance on these metrics, making it easier to quantify the environmental benefits of systems. This capability is particularly valuable as stakeholders demand greater transparency and accountability in sustainability reporting.

Technical innovations driving performance

The engineering advancements in high-temperature fluid coolers have made them a reliable and efficient solution for the modern data centre.

Key features include:

- **Elevated temperature thresholds:** These systems operate efficiently at supply water temperatures of up to 40°C, significantly reducing the need for energy-intensive cooling alternatives.
- **Adaptive controls:** Intelligent systems monitor real-time workloads and adjust cooling output dynamically, allowing optimal energy use without sacrificing thermal stability.
- **Scalable design:** Modular configurations enable operators to expand cooling capacity as needed, making these systems adaptable to growing facilities.

These technical innovations enhance performance and align with the demands of high-density computing environments. As workloads become more intensive and diverse, the ability to scale and adapt will be important in maintaining operational efficiency. By reducing partial Power Usage Effectiveness (pPUE) in comparison to conventional chillers, energy consumption can be lowered, leading to a decrease in operational costs and carbon footprint.

A look ahead

While the benefits of high-temperature fluid coolers are clear, the path forward is not without challenges. The increasing thermal demands of AI and high-performance computing will continue to push cooling systems to their limits. Operators will need to balance immediate performance requirements with long-term adaptability, this means that systems remain viable as technology evolves.

Regulatory frameworks are also likely to grow more complex, requiring ongoing innovation in cooling technology. Data centres will need to stay ahead of these developments, investing in solutions that can adapt to changing standards while maintaining operational efficiency.

On the opportunity side, advancements in materials science and control systems could further enhance the efficiency and sustainability of high-temperature fluid coolers. Integrating these systems with emerging technologies like digital twins and predictive analytics could unlock new levels of performance and reliability, enabling data centres to remain at the forefront of innovation.

Ultimately, the journey towards net zero data centres is defined by complexity and opportunity. High-temperature fluid coolers, as a key component of advanced chilled water systems, offer a practical and impactful solution to the challenges of sustainability, efficiency and compliance.

By enabling heat recovery, reducing energy demand and aligning with stringent environmental regulations, these systems are transforming how data centres approach cooling.

For operators looking to make their facilities ready for the future, investing in high-temperature fluid coolers is a necessary step towards a more sustainable and resilient future.

Why the data centres powering tomorrow's AI will be liquid-cooled

The power demands of generative AI has placed the energy demands of data centres under the spotlight, but this actually offers business leaders an opportunity to embrace innovative solutions and work towards a 'cleaner' approach to AI.

BY NOAM ROSEN, EMEA DIRECTOR, HPC & AI AT LENOVO ISG

THE LAST couple of years has seen a huge surge in excitement around generative AI, with businesses racing to harness its potential. Business leaders who are embracing the technology have been confronted with the challenges of coping with its huge and growing energy requirements.

Training, building and using generative AI models requires vast amounts of energy, thanks in no small part to the fact that the Graphics Processing Units (GPUs) used to power generative AI are highly energy-intensive.

For businesses using AI, there is not only a large increase in power consumption, but also a growing need for dense computational resources. All this works together to create heat.

This has thrown a sharp focus on the rapidly growing energy demands of data centres, and the enormous amount of heat they generate. Business leaders are now facing an urgent need for innovative solutions to manage this heat, and this is where liquid cooling can be a tangible solution.

The problem with air cooling

Energy intensive Graphics Processing Units (GPUs) that power AI platforms require five to 10 times more energy than Central Processing Units (CPUs), because of the larger number of transistors. This is already impacting data centers. There are also new, cost-effective design methodologies incorporating features such as 3D silicon stacking, which allows GPU manufacturers to pack more components into a smaller footprint. This again increases the power density, meaning data centers need more energy, and create more heat.

Another trend running in parallel is a steady fall in TCase (or Case Temperature) in the latest chips. TCase is the maximum safe temperature for the surface of chips such as GPUs. It is a limit set by the manufacturer to ensure the chip will run smoothly and not overheat, or require throttling which impacts performance. On newer chips, T Case is coming down from 90 to 100 degrees Celsius to 70 or 80 degrees, or even lower. This is further driving the demand for new ways to cool GPUs.

As a result of these factors, air cooling is no longer doing the job when it comes to AI. It is not just the power of the components, but the density of those components in the data center. Unless servers become three times bigger than they were before, efficient heat removal is needed. That requires special handling, and liquid cooling will be essential to support the mainstream roll-out of AI.

Turning to liquid

Liquid cooling is growing in popularity. Public research institutions were amongst the first users, because they usually request the latest and greatest in data center tech to drive high performance computing (HPC) and AI. Yet they tend to have fewer fears around the risk of adopting new technology before it is already established in the market.

Enterprise customers are more risk averse. They need to make sure what they deploy will immediately provide return on investment. We are now seeing more and more financial institutions – often conservative due to regulatory requirements – adopt the

technology, alongside the automotive industry.

The latter are big users of HPC systems to develop new cars, and now also the service providers in colocation data centers. Generative AI has huge power requirements that most enterprises cannot fulfil within their premises, so they need to go to a colocation data center, to service providers that can deliver those computational resources. Those service providers are now transitioning to new GPU architectures, and to liquid cooling. If they deploy liquid cooling, they can be much more efficient in their operations.

Across the data centre

Liquid cooling delivers results both within individual servers and in the larger data centers. By transitioning from a server with fans to a server with liquid cooling, businesses can make significant reductions when it comes to energy consumption. But this is only at device level, whereas perimeter cooling - removing heat from the data center – requires more energy to cool and remove the heat. That can mean only two thirds of the energy that the data center is using is going towards computing, the task the data center is designed to do. The rest is used to keep the data center cool.

Power usage effectiveness (PUE) is a measurement of how efficient data centers are. You take the power required to run the whole data center, including the cooling systems, divided by the power requirements of the IT equipment. With data centers that are optimised by liquid, some of them are doing PUE of 1.1, and some even 1.04, which means a very small amount of marginal energy. That's before we even consider the opportunity to take this hot liquid or water coming out of the racks, and reuse that heat to do something useful, such as heating the building in the winter, which we see some customers doing today.

Density is also very important. Liquid cooling allows us to pack a lot of equipment in a high rack density. With liquid cooling, we can populate those racks and use less data center space overall, less real estate, which is going to be very important for AI.

Meeting the challenge

Liquid cooled systems will become an essential solution for businesses dealing with the growing energy demands of generative AI. Today, liquid cooled systems will help business leaders optimise their energy efficiency, but the technology will also enable data centres to deal with the ever-growing number of GPUs required to power future advancements in AI.

Air cooling has simply become inadequate in the age of generative AI. The power demands of generative AI has placed the energy demands of data centres under the spotlight, but this actually offers business leaders an opportunity to embrace innovative solutions and work towards a 'cleaner' approach to AI.

Should you be using power or water to cool your data centre?

Whilst there isn't a simple solution to power vs water-based cooling, understanding how the data centre fits into the wider ecosystem holistically will ensure the correct choice.

BY ANDREW LIVESEY, ASSOCIATE DIRECTOR, CUNDALL

I'D MAKE a bet you've used your phone today. Maybe it was messaging someone. It could be doing some online banking. Likely, you scrolled on social media.

No matter what you did, a data centre was definitely involved. These buildings are so important to daily life that the UK deemed them mission-critical infrastructure so that they are better supported in case of an outage.

To keep these buildings operational requires large quantities of energy and water. For data centre operators who are also trying to reach sustainability targets, their design is vital.
One key aspect of data centre design is the cooling system. The equipment in data centres runs hot, so efficient cooling of the equipment is a primary aspect of operation. It is up to mechanical engineers to select the system that will best suit the data centre. They are typically power-based or water-based, and each comes with advantages and setbacks.

Data centres are in every country, and local climates affect facility's design in different ways. Some countries have a scarcity of water, whereas others could rely on high carbon power grids. When bringing reference designs to different countries, choosing between water or power systems is rarely a one-size-fits-all solution.

Water usage

In the past decade, power usage efficiency (PUE) has been the primary metric for measuring a data centre's sustainability impact. Driving this figure down is the goal, and so many operators use water-based systems. These systems also have two categories, those being direct-air or indirect-air.

In a direct-air system, air from outside the building is brought into the data centre through an evaporative cooling humidifier. The water evaporates as air passes over the humidifier, cooling the incoming air. This system also uses air-side economisers to circulate warm air, which increases efficiency.

The indirect-air system operates similarly. However, in this, the air stream in the data hall is passed over a heat exchanger, and the external air stream is rejected by the atmosphere. This means that the external air does not enter the data hall. To improve the performance of the heat exchanger, water is run over it. Doing this means you don't need to run chiller compressors in the system unless it is in a particularly hot climate. This minimises power consumption and reduces PUE whilst maximising power for the IT systems.

However, the water use of data centres has been a focal point for many when discussing the sustainable impact of data centres. This water could instead be used for residential or agricultural purposes. Additionally, the growth of AI has necessitated liquid cooling to be used in data centres due to the higher power of the data racks. For liquid cooling, a hydronic system is used. This passes liquid through the racks to a series of chillers. This provides a greater cooling density but increases the power consumption of the data centre, in turn raising its PUE. That being said, when water is scarce or clean electricity isn't available, it can be a suitable solution for non-AI data centres.

Location matters

From one country to another, conditions vary greatly. In California, water is scarce, and a drought can heavily impact data centres. Water-based cooling is oftentimes not an option. However, they do have access to relatively clean and reliable power so that they can use power-based cooling instead. On the other hand, a region like Poland operates on the converse. Their power grid relies on fossil fuels, so using them would increase their carbon footprint. The most sustainable cooling solutions rely on the conditions in which the data centre will operate.

Waste heat recovery or using river or sea water are relatively new alternatives for cooling. Waste heat recovery involves capturing the heat generated by the servers, where it is redirected to facilities like district heating networks, making what was once waste a valuable resource. We worked on a hyperscale data centre in Denmark that used heat pumps providing around 165,000MWh of heat per year when the facility was fully built out.

Current sustainability metrics are inaccurate

PUE or WUE – water usage effectiveness – have been traditionally used to measure the sustainability of data centres. PUE measures the total energy consumed ratios against the energy used by the IT equipment. The trouble is that this doesn't reflect the true environmental impact of the data centre. For instance, when water plays a larger role in cooling systems. Meanwhile, WUE measures the amount of water consumed against the IT power. Again, though, WUE can vary depending on the cooling system. Also, it doesn't consider the water used when generating power. This means that it's tough to compare data centre to data centre.

Recently, carbon usage effectiveness (CUE) has been put forward as a better solution. This metric can be best to find what the best cooling solution is for the data centre. This is because it can measure the benefits of reducing power and water against the potential carbon output. To make this most accurate, though, the carbon intensity of the power and water production needs to be included in the calculation.

Consider your location

The number and scale of data centres being built is massive. To speed up their delivery, operators use reference designs using a data centre template and bring it to different regions. This means designers can edit the design rather than start from a blank page. Whilst this does mean that not every data centre needs to be designed from scratch, it doesn't reduce the need for engineering and design teams. When bringing a reference design to a new location, the climate and local regulations must be brought into the design process. These will heavily impact the best cooling solution. Whilst there isn't a simple solution to power vs water-based cooling, understanding how the data centre fits into the wider ecosystem holistically will ensure the correct choice.

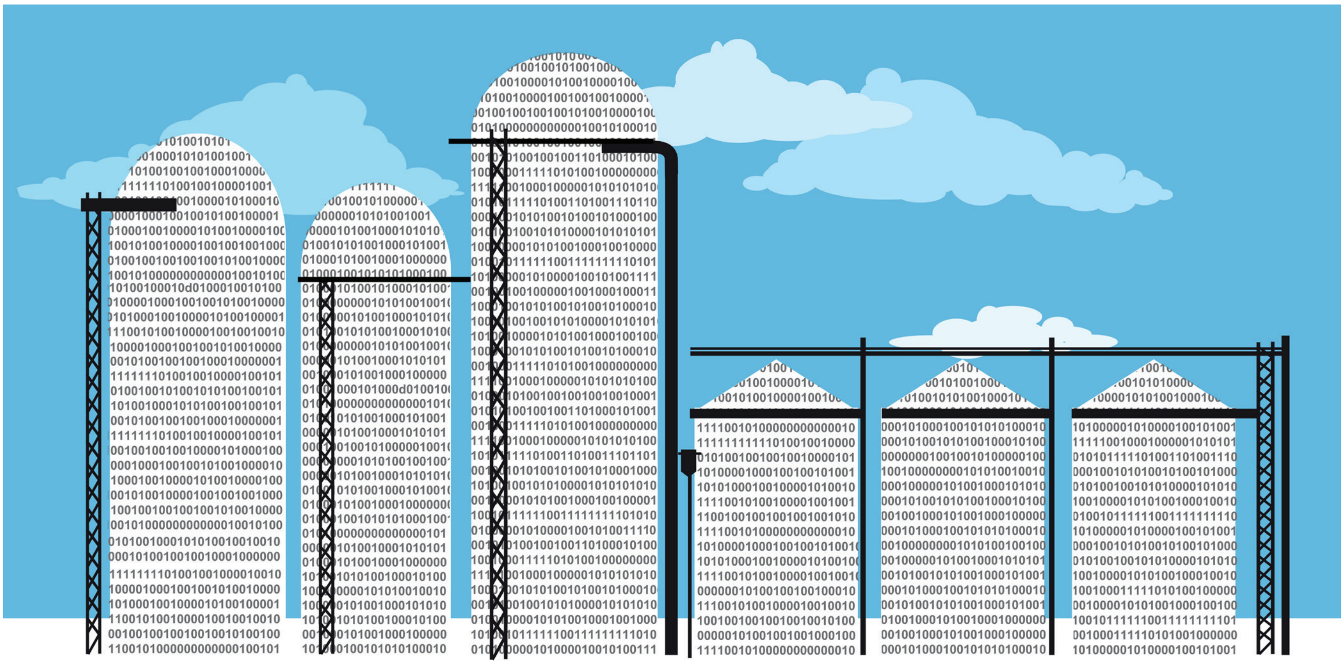

From silos to synergy

The kaleidoscope effect of data centre digital twins

BY AITOR ZABALEGUI, SENIOR PRINCIPAL APPLICATION ENGINEER AT CADENCE

ACCORDING to Salesforce, 81% of IT decision-makers report that data silos are the biggest barrier preventing their organizations from meeting digital transformation goals. Data centers, the backbone of our digital world, are no strangers to this challenge either.

These facilities face mounting pressure to improve efficiency, manage growing demands, and achieve sustainability targets. Yet, many remain stuck in reactive modes, hindered by poor data quality and siloed processes, addressing issues only as they arise rather than proactively identifying and resolving them.

This slow and fragmented approach is no longer sustainable in today's fast-paced environment. Modern data centers must be agile and capable of making informed decisions to meet the growing demands of AI. Enter digital twins. Digital twins are virtual replicas of the physical facility, serving as the connective tissue for varied and siloed data. By consolidating diverse data points into a unified, actionable model, digital twins empower data center operators to act quickly and decisively via predictive insights.

Turning silos into answers

Think of a digital twin like a kaleidoscope. Data, when viewed individually in its respective silos, is often flat and uninspiring. There is insufficient information to see the full picture or make decisions confidently. However, when data is layered into a single cohesive view – via application programming interfaces (APIs) and routine file transfers – it reveals depth, dynamic interplay, and the visibility necessary to illuminate the path forward. Digital twins enable data center operators to see beyond the constraints of their silos, empowering them to mitigate risks, address inefficiencies, and plan proactively.

Many data center operators – including hyperscalers – still rely heavily on spreadsheets, primarily because spreadsheet data is more flexibly structured, simpler to use, and cheaper than using a data center infrastructure management (DCIM) tool. However, the downsides, such as inconsistently formatted data entry, overly complex macros, and lack of interoperability and automation, are becoming harder to ignore. In other cases, outdated tools and incomplete datasets leave teams struggling with missing information at critical moments. For instance, discrepancies in power usage records between planned and actual can create confusion, drive inefficiencies, and even lead to operational risks.

What's more, siloed toolsets and spreadsheet data offer no automatable mechanisms to spot mismatches or gaps, which rear their heads at the most inconvenient times. For example, a system might show that a cabinet is energized, but there's no record of how it is connected to power, raising questions about whether it's safe to add more devices. In another instance, an alarm might go off for a cabinet that is supposed to be empty and not drawing any power, leaving teams unsure if the issue is a faulty sensor or a mislabelled cabinet. Digital twin technology helps solve these problems by providing visibility and confidence that can only come from layering an organization's toolsets, workflows, and data into a unified model.

Digital twins change the paradigm

Digital twins not only highlight data gaps and discrepancies but also give operators the power to resolve them. This process ensures that problems can be addressed proactively rather than reacted to after they have already driven inefficiencies.

Moreover, digital twins excel in contextualizing data within the broader operational landscape of a data center. Integrating data streams from power, cooling, IT, and other systems creates a new, holistic data stream that uncovers insights that would otherwise remain hidden when viewed in isolation. This is especially true when managing airflow, which is invisible by nature yet is revealed by digital twins through computational fluid dynamics (CFD) simulation. These insights are further enhanced through visualization in a 3D model, offering an intuitive and familiar perspective that goes beyond traditional timeline charts, enabling more strategic decision-making.

Ready to embrace the change?

For data centers ready to benefit from digital twins, the journey starts with enhancing data collection and collaboration. Instead of requiring teams to abandon familiar workflows or tools, digital twins seamlessly integrate with existing systems, even leveraging legacy spreadsheets to minimize disruption. As facilities experience the benefits of better insights and reduced manual work, transitioning to a unified data environment becomes a natural progression.

With AI, edge computing, and high-density workloads driving demand, data centers need more than incremental improvements – they need a fundamental shift in how they operate. Digital twins provide this by offering a holistic view of the entire ecosystem, turning "boring" data into actionable insights that deliver real value. Armed with digital twins and improved data quality, data centers can break out of reactive cycles, unlock new efficiencies, and establish a future-ready foundation to tackle tomorrow's challenges.

With AI, edge computing, and high-density workloads driving demand, data centers need more than incremental improvements – they need a fundamental shift in how they operate

The AI Era: redesigning data centres for scale and speed

AI's insatiable thirst for computing resources is transforming infrastructure, and the industry is grappling with how to meet the power, scalability, and efficiency demands. This has spurred an influx of investment to reconfigure data center architectures to address these and other technology requirements.

BY MARIE HATTAR, SVP AT KEYSIGHT

THE CRUX of the issue is that creating the intelligence requires immense computational capacity. With AI complexity increasing by orders of magnitude each year, data centers need to scale rapidly. To provide perspective, demand is growing so quickly that by 2027, AI workloads will consume more energy than Argentina does.

One size doesn't fit all

AI is redefining the architecture of every type of data center: hyperscalers, onsite, colocation, and edge. Most of the attention to date has focused on the hyperscalers arms race. The exponential demand for computational resources is creating AI clusters with sites exceeding 1GW of capacity. McKinsey predicts that by 2030, over 60% of AI workloads in Europe and the US will be hosted on hyperscale infrastructures.

Hyperscale to edge: The architectural spectrum

Data centers must be able to support AI workloads like training large language models (LLMs). This involves overhauling facilities' design and architecture. Power capacity must increase to 200-300kW per rack to support intensive compute, and enhanced cooling solutions that are required at these densities.

Specialized hardware like GPUs and TPUs must be integrated along with expanding storage systems to manage the vast volume of data. Disaggregated architectures are

being deployed so hardware can be managed and scaled independently, enabling different workloads to utilize resources efficiently. Network architectures require updates to handle AI traffic patterns, or AI clusters may become digital gridlocks—processing powerhouses paralyzed by data bottlenecks.

In addition to hyperscale facilities, AI is driving demand for decentralized infrastructure to support data processing locally. This requires centers designed for edge workloads - high performance within a smaller physical footprint and lower energy consumption. By 2030, as more processing shifts to the edge, that market is expected to exceed $160 billion.

This growth is driven by the need to support real-time processing closer to end users for applications like autonomous driving, where faster decision-making is paramount. This approach reduces latency and supports the hyper-connected world fueled by IoT and 5G technologies.

As AI adoption matures, inference workloads are growing at a significantly faster rate than training models. Infrastructure needs to account for this shift from training to inference, which DeepSeek R1 and OpenAI v3 rely on.

These reasoning systems utilize a trained model that evaluates live data to make a prediction or solve a task efficiently.

Connected devices at the edge will generate much of the data. Therefore, facilities need enough scale to support low latency networks with flexible resource allocation. This will allow them to account for the unpredictable spikes in demand for inference.

Scaling for and with AI

Paradoxically, AI is the problem and the solution. Intelligence is vital to solving scaling challenges and ensuring efficient operations. AI can modernize data centers in numerous ways including:

Improving energy efficiency is essential for sustainable operations. Deploying AI can automatically adjust cooling systems and server workloads to meet spikes in demand. Implementing

intelligent energy-saving techniques helps minimize waste and operational costs while maintaining performance levels; Google reduced energy cooling by 40% in its data centers.

Predictive maintenance uses machine learning to anticipate problems before they occur. This minimizes downtime and helps extend infrastructure life.

With the size and costs involved in scaling, the ability to proactively schedule repairs and updates to optimize resource utilization has a material impact.

Digital twins augmented with AI create dynamic models to test and validate components and systems. These solutions can be used to ensure that complex data centers are robust, resilient, and able to support future demands. The AI algorithm analyzes historical data on performance and environmental conditions, providing insights to optimize operations.

The solutions can use AI workloads to emulate network performance to find and fix potential bottlenecks. Advanced test and simulation tools are vital parts of the technology stack needed to create scalable, efficient, and reliable infrastructure.
AI will speed the road towards fully autonomous, intelligent data centers that handle nearly all operations—from monitoring to maintenance to networking to energy management and security—with minimal human input.

Future Proofing AI Infrastructure
As AI matures, data centers must accommodate increasingly complex workloads. Operators are clamoring to scale their infrastructure in a sustainable manner to support the demands without sacrificing performance or reliability.

With much of the AI roadmap still hazy, creating flexible and resilient infrastructure that can easily adapt is vital.

The ability to balance hyperscale muscle with edge agility orchestrated by AI systems will distinguish the winners from the losers.

Those providers that embrace this reality will thrive in the AI revolution, while others will power down.

Rethinking Sustainability

Time to look at the bigger picture and avoid simplistic thinking?

BY EMMA DENNARD, VP NORTHERN EUROPE, OVHCLOUD

LITHIUM is a difficult mineral to extract. Extracting it is primarily done through either evaporation from salty underwater pools, or by traditionally blast-mining spodumene, a lithium-rich mineral.

According to some sources, extracting a tonne of lithium can consume up to half a million litres of water and generates fifteen tonnes of CO_2.

However, lithium is at the heart of many sustainability drives; an EV battery uses around eight kilograms of it. But once it's there, it greatly supports our drive to become greener – for example, EVs put out about 60% of the emissions that traditional petrol vehicles do.

The point of the exercise is to show that we need to look at the bigger picture and avoid simplistic thinking. Progress always has a cost, and we shouldn't flinch from looking at the broadest possible context to make sure that our actions don't have unintended consequences.

Seizing the opportunity to rethink our industry

The cloud industry has developed some bad habits, and it's clear that in some cases, we've built on bad foundations. There's a reason why a certain brand of electric car is popular: the designers went back to the drawing board and thought about what a truly modern car should do and how it should do it, retaining the good without being afraid to throw out the bad.

In our industry, for example, there is an unhealthy fascination with using air in datacentres. Compared to water, air is difficult to direct, and the volumetric heat capacity of water is thousands of times better than that of air. Although air is useful in removing heat from some areas, water is far more efficient when cooling very specific and well-defined areas, such as CPUs and GPUs. According to studies, an average air-cooled datacentre has a PuE of around 1.55, compared to 1.2 for Direct Liquid to Chip (DLC) cooling.

And as we know, liquid cooling can be taken further by immersing the entire server in liquid or even combining this with DLC cooling to ensure that the hottest parts of the server get the attention they need.

Clearly, this does make maintenance harder, but liquid cooling at scale has

When whole components can't be re-used, it's possible to re-use certain metals. For example, using electrolysis, you can separate copper and other metals from even the most complex and tricky part of the server, the motherboard. This copper can be re-sold and re-used, turning waste into revenue

been deployed by many organisations – including ours – for a number of years.

Thinking bigger

However, reinventing how we think about technology has to go far beyond manufacturing and operations – it's also crucial to think about what happens at end-of-life, which brings us back to minerals.

When servers reach end-of-life, it's rarely the case that all the components need replacing at the same time. The unit might be nearing failure on certain benchmarks, but that doesn't mean parts of it can't be re-used. Some parts – hard drives, for example – may need to be securely disposed of in specific cases to comply with data protection regulations, but other parts can be removed and re-used in lower-spec ranges.

And when whole components can't be re-used, it's possible to re-use certain metals. For example, using electrolysis, you can separate copper and other metals from even the most complex and tricky part of the server, the motherboard. This copper can be re-sold and re-used, turning waste into revenue.

However, there's also a geo-political component to this. Much of our copper comes from Poland and Germany, but around 70% of the EU's cobalt, used in batteries, and almost a third (29%) of the EU's nickel (nickel-plated copper is common in heatsinks) used to come from Russia (European Central Bank, 2023).

Magnesium, to give another example, is frequently used in laptops and parts with complex casings because it's both lightweight and strong – but in 2023, 97% of the EU's magnesium was processed in China. Furthermore, there's also a need to take stock and re-examine the materials we use. For example, could we use aluminium rather than copper for cooling components? With some work and in some cases, yes.

Aluminium is about four times cheaper than copper, because it's over a thousand times more common. It does take more energy to produce, but in the long run, it emits less carbon and tends to be less harmful to marine environments when disposed of poorly, for example.

On the other hand, it's worth noting that we do have to be a bit discerning. Many of the new towns in the UK were created with aluminium rather than copper telephone wires because of post-war shortages.

Aluminium only conducts sixty percent as well as copper, and as a result, areas like Milton Keynes experienced atrocious broadband speeds for some years, until they were re-wired with copper and fibre optic cables later on.

Getting used to complexity

Thinking broadly can be taxing, but it's worthwhile. After all, as we've shown, using more efficient technologies can directly cut power bills, and recycling rare metals can be a source of revenue.

There's no doubt that our industry is complex, and sometimes it's easy to only see the challenges.

However, it is becoming increasingly clear how important it is to understand every link in the chain so that we can tackle the problems honestly, seize the opportunities when they arise, and make sure that we take even the smallest steps towards being a better, cleaner industry.

Because as the saying goes, a journey of a thousand miles still begins with the first step.

DCS AWARDS 2025

Celebrating 15 Years of Industry Excellence

On 22 May 2025, Angel Business Communications proudly hosted the 15th annual DCS Awards at the Leonardo Royal Hotel London St Paul's. The evening brought together over 300 senior professionals from across the data centre industry for a night of recognition, networking, and celebration, honouring those who continue to drive innovation and excellence in the sector.

This year's awards reflected a particularly buoyant period for the data centre industry, as it undergoes rapid transformation driven by surging demand for AI, cloud computing, and digital services across all sectors. The DCS Awards recognised outstanding achievements in sustainability, facilities innovation, ICT advancements, and customer service.

DCS PROJECT AWARDS

NEW BUILD (GREENFIELD OR BROWNFIELD) PROJECT OF THE YEAR
Sponsored by H2H Communications
WINNER – Vantage Data Centers

DATA CENTRE SUSTAINABILITY PROJECT OF THE YEAR
Sponsored by Centiel
WINNER – EkkoSense and CBRE

DATA CENTRE CONSOLIDATION / UPGRADE PROJECT OF THE YEAR
WINNER – Schneider Electric and APT

CLOUD PROJECT OF THE YEAR
Sponsored by Huawei Digital Power Technologies Co., Ltd
WINNER – EnerSys

DCS INNOVATION AWARDS

DATA CENTRE FACILITIES INNOVATION AWARDS

DATA CENTRE POWER INNOVATION OF THE YEAR
WINNER – Huawei Digital Power Technologies Co., Ltd.

DATA CENTRE POWER DISTRIBUTION UNIT OF THE YEAR
WINNER – Panduit

DATA CENTRE COOLING INNOVATION OF THE YEAR
WINNER – VERTIV

DATA CENTRE SUSTAINABILITY INNOVATION OF THE YEAR
WINNER – Clean Energy Capital

DATA CENTRE INTELLIGENT AUTOMATION / MANAGEMENT INNOVATION OF THE YEAR
WINNER – RiT Tech

DATA CENTRE SAFETY & SECURITY INNOVATION OF THE YEAR
WINNER – Pure Data Centres Group

DCIM / DC MANAGEMENT INNOVATION OF THE YEAR
WINNER – RiT Tech

EDGE INNOVATION OF THE YEAR
Sponsored by Pulsant
WINNER – StorMagic

DATA CENTRE ICT INNOVATION AWARDS

DATA CENTRE ICT STORAGE INNOVATION OF THE YEAR
WINNER – Geyser Data

DATA CENTRE ICT SECURITY INNOVATION OF THE YEAR
WINNER – Hornetsecurity

DATA CENTRE ICT NETWORKING INNOVATION OF THE YEAR
WINNER – Bridgeworks

DATA CENTRE ICT AUTOMATION / ORCHESTRATION INNOVATION OF THE YEAR
WINNER – Türk Telekom

DATA CENTRE ANALYTICS / OBSERVABILITY INNOVATION OF THE YEAR
WINNER – Gapit Nordics AS

COLOCATION EXCELLENCE AWARDS

MANAGED SERVICES / COLOCATION PROJECT OF THE YEAR
WINNER – atNorth

BEST COLOCATION PROVIDER SUSTAINABILITY PROJECT OF THE YEAR
WINNER – nLighten

BEST COLOCATION PROVIDER SUSTAINABILITY INNOVATION OF THE YEAR
Sponsored by Huawei Digital Power Technologies Co., Ltd
WINNER – Therma Loop AB

DATA CENTRE HOSTING/ COLOCATION SUPPLIER OF THE YEAR
WINNER – Conapto

DCS COMPANY AWARDS

DATA CENTRE MANAGED SERVICES VENDOR OF THE YEAR
WINNER – OryxAlign

DATA CENTRE FACILITIES VENDOR OF THE YEAR
WINNER – Huawei Digital Power Technologies Co., Ltd.

DATA CENTRE ICT SYSTEMS VENDOR OF THE YEAR
WINNER – Hubnetix

OUTSTANDING CONTRIBUTION TO SUSTAINABILITY AND EFFICIENCY AWARD
Sponsored by OryxAlign
WINNER – Sudlows

CUSTOMER-CENTRIC SERVICE EXCELLENCE AWARD
WINNER – Telehouse

EXCELLENCE IN DATA CENTRE SERVICES AWARD
WINNER – Critical Facilities Solutions

OUTSTANDING CONTRIBUTION TO DIVERSITY, EQUALITY AND SOCIAL RESPONSIBILITY AWARD
WINNER – Kao Data

OUTSTANDING CONTRIBUTION TO SKILLS AND TRAINING AWARD
WINNER – CBRE

BEST OVERALL DATA CENTRE OF THE YEAR
Sponsored by OryxAlign
WINNER – Green Mountain

DCS INDIVIDUAL AWARDS

DATA CENTRE INDUSTRY CONTRIBUTION OF THE YEAR
WINNER – Andy Hirst - Sudlows

EMERGING TALENT AWARD
WINNER – Dena Rahman - Pure Data Centres

EDITOR'S CHOICE AWARD
Sustained Innovation in the Data Centre
centiel continuous power availability
Schneider Electric

The DCS Awards team would like to thank all those who participated this year, especially our sponsors including: Headline sponsors Huawei, Silver sponsors Hyperview, and Category sponsors included Centiel, H2H Communications, Oryx Align Ltd, and Pulsant. We'd also like to thank our supporters The Data Centre Alliance.

Failure is not an option as AI downtime will be the most expensive ever

AI scaling challenges for MEP equipment supply and operation – and how to cope with them

BY AARON BALL, VICE PRESIDENT EMEA, SERVERFARM

DATA CENTRES are national critical infrastructure, and as was recently demonstrated when London Heathrow Airport was forced to close through loss of power, when critical facilities stop working the impact is costly, lengthy and expensive.

What words best describe data center design and operations challenges in 2025? Complexity? Scale? Density? Liquid?

That AI, power densification, and liquid cooling have created new pressures for data center design and operation engineering teams is no longer news to anyone with any kind of responsibility for 'keeping the lights on' in a data hall. That changes are happening 'in flight' as supply chains become stretched and waiting times for new equipment lengthen provides an interesting context.

The bar on the risks associated with equipment failure has been raised. Downtime is becoming more expensive than ever. The question is: Are the frequency of M+E equipment failures rising at a time when service criticality is growing?

These new layers of complexity have hit just when capacity scale out is the new normal and at a time when speed to market is a key customer requirement. It is enough to cause sleepless nights.

Don't Panic

Readers of a certain age will recognise the words "Don't Panic" from the iconic Hitchhiker's Guide to the Galaxy series of books. While we are not dealing with the end of the Universe, it would be foolish to pretend that there may be more to supply chain issues than simply sourcing equipment.

So, while we are not panicking we know we are dealing with supply chain disruptions which can't be ignored. According to Datamation Global supply chain issues have extended lead times for critical components. Electrical equipment like switchgear and generators, which previously had lead times of up to six months, now may take up to 18 months to two years to arrive.

There are also challenges in construction material delays. A 2023 survey reported by Statista indicated that the majority of respondents experienced an increase in data center construction material lead times, with respondents reporting delays of more than eight weeks on average.
It is no surprise the demand for MEP equipment for critical services has never been greater.

Huge growth in the data center industry, alongside mass electrification across every sector from automotive to marine transport, to manufacturing mean power chain equipment makers have full order books. This has come in the aftermath of the huge impact of the Covid pandemic on global supply chains.

Increased complexity and incident responses to multiple failures?
As digitalisation accelerates across every aspect of life, data centre operators have become the infrastructure foundation for the world's critical services.

Colocation service providers such as Serverfarm want to invest in equipment of high build quality, that operates well in the field, needs minimal servicing and does not fail.

Of course, no-one likes to talk about equipment failures. But failures are a reality. And given the pressures being felt by manufacturers it is a reasonable question to ask if we are going to see more equipment failures in the future.

At this time of surging demand the challenge for manufacturers is to ensure quality does not suffer. But investing the quality equipment is just part of the solution. Effective day to day operations is knowing what to do when serious incidents occur – even where a cascading series of events might see multiple concurrent equipment fall over.

For example, everyone who uses a colocation provider has asked: "What happens if you lose the grid?" In normal circumstances losing the grid means transferring to a second feed.

But how many follow up with the question - what happens if there is an equipment failure and the second feed is offline? And what then happens if one of the gensets fails to start up? How much pressure can the UPS and battery equipment take? And what happens if a failure happens here too?

Is such a series of events unimaginable? Not if you are an engineer.
Were such events to occur in sequence is exactly the time when standards, consistency, processes and professional expertise can prove the difference between 'incident' and 'catastrophe.'

In a well run data center operational expertise includes the on site teams having forensic insight into the design and construction of every aspect of the facility. The best way to achieve this is through those teams participating at the design and construction stage of the data center.

This level of collaboration means that operations can see 'how this facility is actually going to run' is understood even before the first GPU server is wheeled into a hall and onto a rack. This must be backed by demonstrating the highest operational standards.

Having proven incident response processes (which are globally consistent, drilled and practiced) and having the right expertise on site could make the difference between maintaining uptime and total shutdown. Furthermore, being open, transparent and engaging with customers provides reassurances and communicates: 'This is how we run things.'

Future tense

The last 10 years have seen major changes in how data centers are operated. The next 10 years will see greater and fundamental changes at every layer of the stack.

In terms of uptime, emergency incident response and planned downtime these changes will come in many forms. More resilience will be baked in at different software layers, from the application to system software.

At the platform layer Cloud-based failover and disaster recovery strategies will provide business continuity by leveraging multi-cloud and hybrid cloud environments, with automated disaster recovery (DR) systems activated instantly during outages. AI-powered failure analysis will prevent disruptions and enhance recovery strategies. Redundant hardware combined with AI-driven infrastructure ensures seamless reconfiguration or replacement components. Edge computing and distributed architectures minimize the impact of failures by decentralizing processing, ensuring that workloads can continue.

As these innovations continue to evolve, data centers will become more resilient, efficient, and capable of handling unexpected failures with minimal downtime.

But those running data centers for AI and cloud providers today must not become over distracted by outside noise and new shiny things.

At an operations level we can embrace new methods and technologies without falling into the trap of becoming over focused on future solutions and missing today's SLAs.

Today's reality

Supply chains are not simply about "How quickly can I get it?" If failure rates creep up then speed of supply is not the answer. A focus on build quality that provides improved hardware resilience contributes to our ability to provide service reliability.

At Serverfarm we know that we live in a design and operations world of extended lead times for critical equipment, and we must plan carefully to meet project timelines. Proactive measures such as phased delivery strategies are essential to mitigate potential disruptions and ensure seamless operations.

We opened this article with a list of the challenges we face in design and operation of modern data centers. For us to serve our customers we want to invest in equipment of the highest build quality, with long MTTM and even longer MTTR numbers.

When we invest in equipment we calculate TCO not simply through a capex and opex lens but also through a quality of service lens for our customers. This is our definition of not failing.

Super-TCXO enables new architecture for greater AI workload efficiency in data centres

AI workload efficiency takes centre stage.

BY SITIME

ARTIFICIAL INTELLIGENCE (AI) is accelerating advancements unlike any technology before. AI can analyze massive datasets and deliver answers in fractions of a second. It's driving progress, everything from immensely capable virtual assistants to new medical breakthroughs and an even better understanding of the universe's origins.

Along with these advancements, datacenters must contend with the challenges of increased energy consumption and rapidly rising costs. To help address these challenges, AI workload efficiency is taking center stage.

The new SiTime SiT5977 Super-TCXO® helps system engineers better optimize AI workload efficiency by enabling new architectures in smart network interface cards (Smart NICs), acceleration cards, switches and compute node applications. The new SiT5977 Super-TCXO is unique in the timing industry—enabling AI workload efficiency with 3X better synchronization and 800G network connectivity in a 4X smaller footprint.

How precision timing can enable greater AI workload efficiency

By improving AI workload efficiency, datacenters can lower total cost of ownership (TCO) and improve energy profiles and carbon footprints. "Improving AI workload efficiency to reduce energy consumption and carbon emissions is an industry-wide challenge. Precision timing is one of the approaches to help solve this problem," said Dave Altavilla, president and principal analyst at HotTech Vision & Analysis.

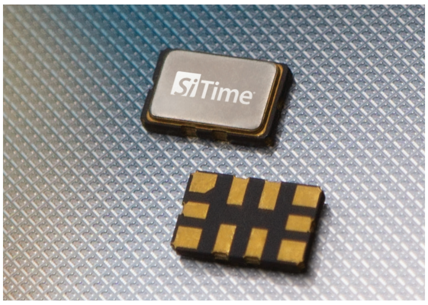

The new SiT5977 Super-TCXO is unique in the timing industry–enabling AI workload efficiency with 3X better synchronization and 800G network connectivity in a 4X smaller footprint.

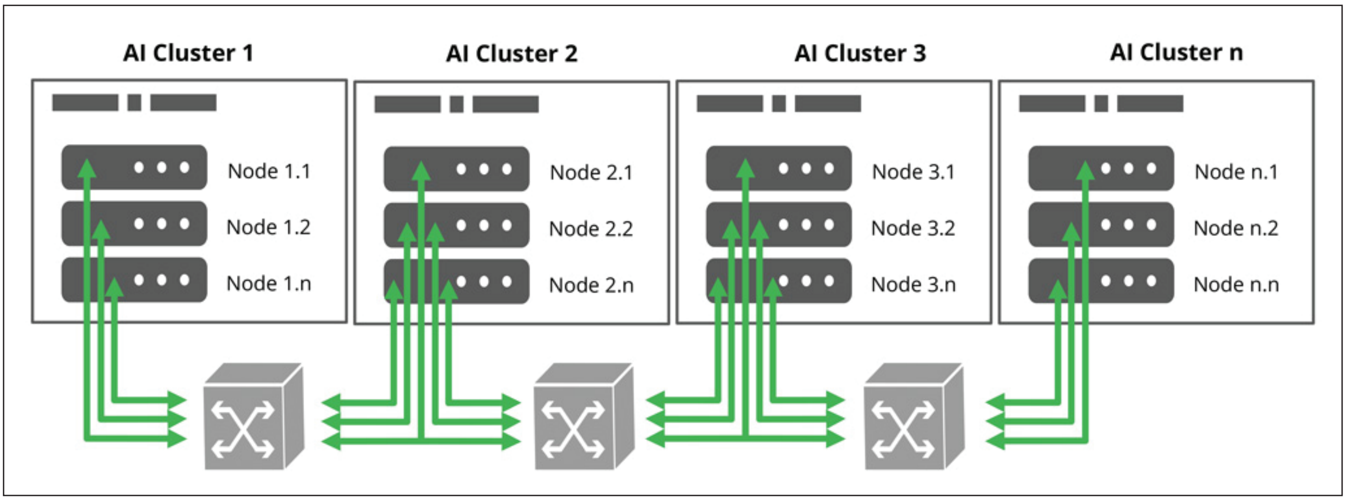

AI workloads are accelerated by connecting nodes and accelerators directly. Faster and higher utilized bandwidth reduces the time to complete a workload.

"SiTime is the only semiconductor company fully dedicated to developing innovative timing solutions required for the complex scaling of today's AI datacenters."

Large AI clusters contain tens of thousands of accelerators, and AI workloads are divided among all the accelerators. The maximum utilization of the cluster depends on both end-to-end network synchronization and high data bandwidth.

Synchronizing the data transmission with precision timing enables proper orchestration of AI data workloads among the accelerators to ensure maximum efficiency. Precision timing's highly tuned synchronization also enables precise telemetry, allowing insights for operators to eliminate underperforming accelerator nodes. In addition, tighter synchronization allows architects to reduce the guardband between data packets, called the window of uncertainty.

A smaller window of uncertainty reduces the chance of overlaps, therefore reducing the frequency of data re-transmits.

Equally important to synchronization is fully utilized network bandwidth. To ensure that AI workloads are done as quickly and efficiently as possible, it's critical to remove bandwidth bottlenecks and ensure that the high bandwidth can be fully leveraged by the compute node. Together, precision synchronization and higher bandwidth utilization increase the efficiency of the entire datacenter.

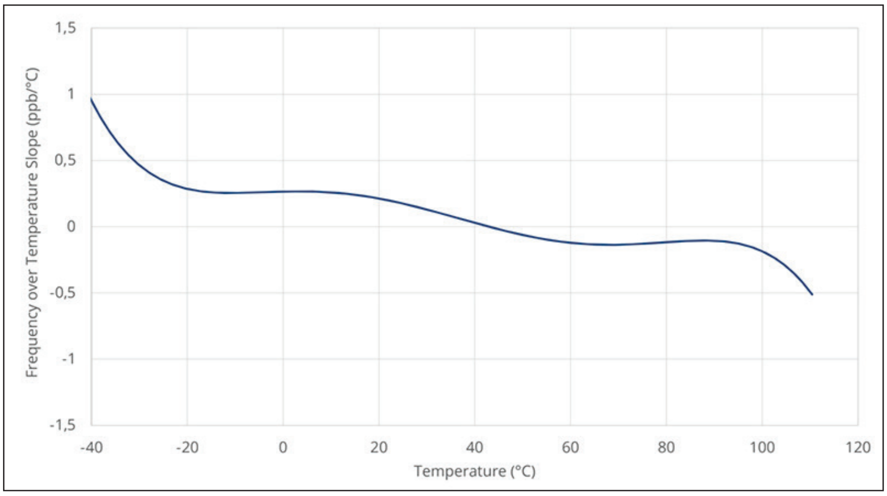

The SiT5977 operates at ±1 ppb/°C stability over temperature slope (dF/dT) to ensure a precise time reference under thermal shock and airflow. This means no performance throttling, fewer dropped packets and smaller windows of uncertainty.

SiT5977 Super-TCXO enables streamlined Ai architecture

The SiT5977 Super-TCXO streamlines AI system architecture. This new chip integrates a low-jitter PLL to support 800G bandwidth with SiTime's ultra-stable DualMEMS™ stability engine.

The SiT5977 further integrates digital frequency control to enable system-level synchronization. It's the only single-chip timing solution that delivers the most resilient performance for AI compute nodes with high bandwidth and network synchronization.

The SiT5977 Super-TCXO performs much better than quartz timing devices which have poor resistance to environmental vibration and airflow. Thermal shock and airflow in the datacenter can disrupt quartz oscillator precision and lead to dropped packets and poor synchronization. The SiT5977 operates at ±1 ppb/°C stability over temperature slope (dF/dT) to ensure a precise time reference under thermal shock and airflow. This means no performance throttling, fewer dropped packets and smaller windows of uncertainty.

In addition to the leading dF/dT of the Si5977, its ultra-low jitter, just 80 fs from 12kHz to 20MHz, is crucial to enabling 800G bandwidths for the most capable AI datacenters. Bottlenecks in bandwidth cause the

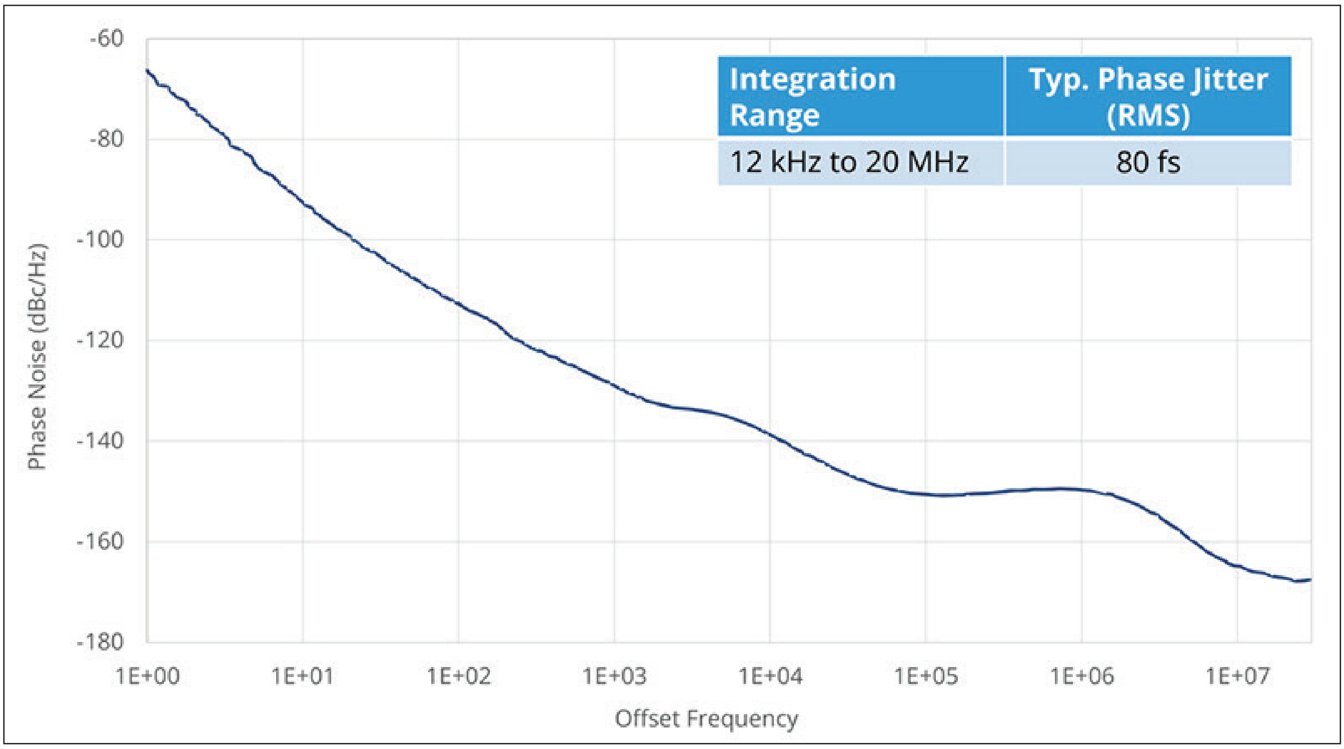

➤ SiT5977's class-leading jitter performance removes data bottlenecks, ensuring that accelerators can run as efficiently as possible.

accelerators to idle while they wait for data, sometimes up to 50% of their run time. High bandwidth in the datacenter infrastructure, including the NIC, is critical to reducing the accelerator idling time. SiT5977's class-leading jitter performance removes data bottlenecks, ensuring that accelerators can run as efficiently as possible.

By integrating the high stability over temperature for nanosecond synchronization, ultra-low phase jitter for 800G bandwidth and high-resolution digital frequency control into a 5 mm x 3.5 mm package, the jitter cleaner used in the traditional quartz architecture can be completely removed. The 4X smaller footprint, compared to the competing architecture, enables larger processors in compact systems. This allows system architects to fully leverage high-speed 800G bandwidth network connectivity and maximize utilization of the AI cluster.

By integrating the high stability over temperature for nanosecond synchronization, ultra-low phase jitter for 800G bandwidth and high-resolution digital frequency control into a 5 mm x 3.5 mm package, the jitter cleaner used in the traditional quartz architecture can be completely removed.

SiT5977 FEATURES

- Environmentally robust with ±1 ppb/°C frequency slope (dF/dT) for optimum performance under airflow, thermal shock.
- Capable of driving 800G and higher links via 80 fs phase jitter and LVDS outputs.
- Enables embedded control loops with precise digital tuning of output frequency (DCTCXO), ±400 ppm pull range and 0.05 ppt (5e-14) resolution via I2C/SPI.
- Eliminates link flaps from quartz timing activity dips or micro jumps.
- Resistant to shock, vibration and board bending.
- Eliminates external LDOs via on-chip voltage regulators. ±100 ppb frequency stability over temperature with DualMEMS ™ architecture.
- 156.25 MHz output frequency enabling high-speed SerDes and 800G links.

Distributed applications, such as AI training and inference, need precision timing to synchronize workloads. The SiT5977 enables a new, more efficient architecture to lower energy costs and optimize bandwidth utilization, helping to unlock the full promise of AI.

Advanced networks for artificial intelligence and machine learning computing

In today's rapidly evolving technological landscape, the demand for artificial intelligence (AI) and machine learning (ML) capabilities is driving significant transformations across industries. As these technologies continue to advance, the infrastructure supporting them must also evolve to meet the increasing demands of computational power and data processing. This is where the groundbreaking white paper, "Advanced Networks for Artificial Intelligence and Machine Learning Computing," becomes an essential resource for industry professionals.

What's Inside:

Energy Consumption: AI data centers consume significantly more power than traditional setups due to the complexity and size of datasets required for training and inference. As models grow, so do their energy requirements, making it essential for data centers to adopt energy-efficient technologies and practices. The document explores various strategies to manage these demands, emphasizing the role of advanced accelerators and energy-efficient hardware.

Cooling Solutions: The high-density environments of AI data centers generate considerable heat, necessitating innovative thermal management techniques. Traditional air-cooling methods are no longer sufficient. Instead, advanced solutions such as direct-to-chip and immersion cooling are essential for maintaining optimal operating conditions.

These methods not only manage heat more effectively but also contribute to reducing overall energy consumption, highlighting the importance of adopting cutting-edge cooling technologies. Network topology: The white paper discusses various topologies, including Clos, torus, and hybrid configurations, each offering unique advantages in terms of data flow efficiency and scalability.

For instance, Clos topologies provide non-blocking, high-bandwidth connectivity, minimizing congestion and supporting efficient data transfer. These topologies are crucial for AI data centers that require rapid scalability and robust performance under increasing demands.

Scalability is a recurring theme throughout the document, as the ability to expand and adapt is vital for future-proofing AI infrastructures. The white paper outlines strategies for building scalable networks, such as adopting a modular approach that allows for seamless integration of new technologies and components. This adaptability ensures that data centers can continue to meet the growing demands of AI as new models and applications emerge.

AFL stands at the forefront of providing the innovative fiber network solutions essential for modern AI data centers. With a strong track record of delivering high-performance, energy-efficient optical fiber solutions, we are uniquely positioned to support the evolving needs of the industry.

For industry stakeholders and decision-makers, downloading the white paper offers a wealth of insights into the complexities and requirements of AI data center operations. It serves as a comprehensive guide to understanding the current landscape and preparing for the future of AI and ML technologies. By leveraging the knowledge shared in this document, data center operators and IT professionals can equip themselves with the tools and strategies necessary to support the next generation of AI innovations.

Download the white paper today and take the first step towards a more efficient, scalable, and innovative AI future.

TECTON II

DESIGN BY
pininfarina

- Best in Class Energy Efficiency
- Modular by Design
- Rapid installation
- Circular design for sustainability
- Easy maintenance
- Best light quality

Technologies that are light-years ahead.
Lasts longer. Sees further.

The ideal solution for Data Centre environments.

DCA Update

BY STEVE HONE, CEO THE DCA – DATA CENTRE ALLIANCE

2026 is moving on at a fast pace, we'll soon be in the summer and thinking about the breaks we have planned. However, focusing on the here and now, there are plenty of Data Centre events planned in June and July. The DCA team is supporting all of the events listed below in various capacities and we look forward to meeting up with our Members and Partners that are attending.

The DCA will be hosting another 10X10 Data Centre Update Briefing on the 30 July in the Wren Suite at St Pauls' Cathedral, this time it's followed by a summer networking reception.

DC EVENTS 2025

Here are just some of the events listed for June and July.

- Data Centre World Frankfurt 2025
- Data Cloud Global Congress, Cannes
- Data Center Expo North America, Santa Clara
- Pan African Data Centres 2025, Johannesburg
- PowerEx Live Conference, Belfry
- DCA 10X10 Data Centre Update Briefing & Summer Networking, London

To see details of these events and the rest of 2025

Data Centre Solutions – Issue 04

Issue 04 of Data Centre Solutions has a number of topics of focus including Power & Cooling, Design & Construction & ESG (Environmental, Social and Governance).
The DCA feature is comprised of articles from DCA Partners and Industry Experts. I'd like to thank all the authors for providing their contributions.

Power & Cooling

- Matthew Thompson, Managing Director Europe at Airsys, explains why liquid cooling solutions must keep developing and evolving.
- Lawrence Leask FInstR, Managing Director at Kaizen Energy Consultancy details of a survey related to large users of refrigeration (DC's included) funded by The Department for Energy Security and Net Zero (DESNZ)

Design & Construction

- Christian Goldsmith, Senior Program Director, Technology at Arcadis delves into the demands of complex, end-date critical construction programs.
- Matt Edgley, Chief Operating Officer at Datum Datacentres shares the positive experiences of building Datums new facility MCR2.

ESG - Environmental, Social and Governance

- Charles Johnston, Sustainable Energy First provides his thoughts on how DC's are leading the way (and have been for some time) by creating a blueprint for sourcing sustainable energy.

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 Hone

Advancing liquid spray cooling

Engineering Insights into Airsys' next generation liquid cooling solution

BY MATTHEW THOMPSON, MANAGING DIRECTOR EUROPE, AIRSYS

LIQUID COOLING TECHNOLOGY is here to stay. Any new development or enhancement in this field must hit the ground running to keep pace with rapid advancements in chip technology. As data centres increasingly contend with the demands of high-performance computing (HPC) and artificial intelligence (AI), the importance of efficient cooling solutions continues to grow. The rapid increase in chip thermal design power or TDP is already rendering some legacy cooling technologies obsolete. In the next 5 years, if customers want their infrastructure investments to remain relevant for the future generations of GPUs and CPUs, forward-thinking is critical. Specifically, solutions must be designed to handle tomorrow's challenges—such as 4 kW+ per-chip heat loads and a fragile global supply chain.

In this context, today's CapEx is really about tomorrow's operational stability. That's why Airsys' focus on high-density capabilities (>4 kW/chip), system standardisation, long-term spare parts availability, and energy recovery directly aligns with what forward-looking customers need.

Airsys, a market-leading cooling solutions provider who constantly strives to stay ahead of the curve and deliver solutions for customer's needs, has developed the next generation of their award-winning liquid spray cooling solution, the LiquidRack. While still utilising forced convection spray technology to achieve a superior coefficient of heat transfer compared to other single-phase technologies, the next-generation LiquidRack features a series of design refinements that further enhance thermal performance and ensure optimal uptime and serviceability, setting the LiquidRack design apart with its unparalleled reliability, low OPEX and ultra-low power usage effectiveness.

While the first generation established the fundamental architecture, the second generation enhances the design with a series of improvements that address the challenges faced by Data Centre Operators regarding practical deployment and uptime, with an emphasis on modularity and standardisation. The first generation LiquidRack housed servers within a single shared drawer, however the new generation architecture segments the system into independent, self-contained server cassettes. This provides several operational advantages. Firstly, it enables individual servers to be serviced or replaced rapidly through quick-release couplings, if needed, enabling servicing to take place away from the white space, in a separate workshop or lay down area. Each server cassette comes pre-installed with its own key components such as low energy pump and heat exchanger.

Additionally, the server cassette design mitigates the risk of cross-contamination or fluid leakage between servers, as each cassette contains its own minimal volume of dielectric fluid. The modular design introduces scalability and compatibility with diverse hardware configurations, accommodating evolving IT requirements with minimal structural changes and is designed for full compatibility within existing data centre infrastructure, requiring no major retrofitting or facility upgrades. This allows for seamless integration into existing rack environments and reduces barriers to adoption, especially in legacy sites or those operating under strict architectural constraints due to its vertical design and low operating weight thanks to the low volume of dielectric fluid required.

This new generation of LiquidRack also includes the integration of an auto-lift system, which allows server cassettes to be safely removed and lowered in a timely manner for server maintenance, then easily returned into position after servicing. This design eliminates the need for manual lifting or third party lifting apparatus, which significantly reduces the effort and risk involved in accessing hardware, therefore improving the efficiency and safety of routine maintenance.

These design shifts ensure optimal uptime and the modular approach not only simplifies maintenance and containment, but also facilitates

rapid replacement, minimising downtime. In cases of component failure, affected server modules can be swiftly swapped with preconfigured spares, ensuring minimal disruption to server operation. This approach aligns with the increasing emphasis on spare part availability in mission-critical environments, where even brief interruptions can result in significant operational or financial losses. The new look LiquidRack builds upon innovations that began with the first generation with the use of 3D-printed spray heads. These spray heads are customised for specific server layouts, allowing tailored cooling performance for every server. The use of 3D printing equipment also facilitates rapid prototyping and localised production, enabling faster iteration and deployment of server-specific solutions.

Another important aspect of the system is its environmental efficiency, as the system also a water usage effectiveness or WUE of 0, and utilises recycled materials in sprayer components. The LiquidRack also typically only requires up to 10 litres of dielectric fluid per server, due to its cassette-based design. This low volume requirement provides improved operational expenditure and a more sustainable solution, whilst also resulting in a low total charged weight, when compared to immersion, that is suitable for retrofitting into existing data centres.

The Airsys solution does not heavily rely on the specific properties of any particular dielectric fluid, giving customers the flexibility to choose their preferred brand based on their priorities—whether environmental impact, sustainability, energy efficiency, or logistics. This approach puts customers in a strong position to align cooling strategy with broader operational and ESG goals.

Furthermore, depending the heat density, the system continues to operate with no need for chilled water plant and instead relies on dry coolers. Heat recovery and reuse remain critical features, gaining increasing relevance as data centres strive toward Net Zero and Carbon Negative sustainability goals. By elevating fluid temperatures within the LiquidRack, Airsys is able to offer 100% heat recovery, with return hot water temperatures of 65°C, which is compatible with district heating and hot water systems.

These features position the LiquidRack as a compelling option for facilities seeking to reduce their environmental footprint while maintaining high-performance computing capabilities.

In summary, this evolution of the LiquidRack is not merely incremental, but part of a broader rethinking of how liquid cooling technologies should evolve to meet the diverse, demanding, and increasingly sustainability-focused needs of modern data centres.

As the industry continues to advance toward higher density and greater energy efficiency, innovations like those embodied in the LiquidRack will play a central role in defining the future of data centre thermal management, and the role data centres can play in the community.

Contact Airsys here

Running the rule over refrigeration

In 2023 The Department for Energy Security and Net Zero (DESNZ) funded a project looking at large users of refrigeration. The aim of the project was to provide an understanding of refrigeration related energy usage and emissions across several sectors including data centres.

BY LAWRENCE LEASK, MD OF KAIZEN ENERGY CONSULTANCY

IN 2023 The Department for Energy Security and Net Zero (DESNZ) funded a project looking at large users of refrigeration. The aim of the project was to provide an understanding of refrigeration related energy usage and emissions across several sectors including data centres.

Kaizen Energy Consultancy in conjunction with the Institute of Refrigeration and Carbon3IT Ltd worked on the surveying different data centres across the UK, the results of all the surveys were presented blind.
Across the 10 sites surveyed we found 19 different types of refrigeration system being used, every site had opportunities to improve efficiency by at least 30%.

At a follow up meeting at DESNZ we were able to compare results with Dr Dermott Cotter FInstR from Star Refrigeration who had completed surveys for the food storage sector. We found identical issues across both sectors; all of these had been highlighted in a previous reports.

The types of refrigeration system can be split into two distinct groups, Group 1 consisted of equipment which used a secondary refrigerant and distributed chilled water. Group 2 consisted of

standard smaller direct expansion (dx) refrigeration system supplying refrigerant to CRAC units located in the halls.

At every site we found dirty fouled air cooled condenser coils, fouled condensers coils will reduce heat transfer, increase condensing temperatures, reduce the rated cooling capacity and drastically reduce system performance.

As condensing temperatures increase two things happen, the absorbed power from the compressor increases (which accounts for 95% of the systems energy consumption) and the cooling capacity is reduced.
The rule of thumb is for each 1°C increase in condensing temperature refrigeration efficiency reduces by 2-4%, it is not uncommon to see a 10°C reduction in condensing temperatures from a full deep clean, this represent a 20%-40% reduction in energy use.
Condenser coils should be cleaned in the opposite direction of airflow (counterflow), this may involve removing condenser fans, getting safe access can

make the task more difficult and may involve safety platform.

Cleaning condenser coils in the wrong direction will bury dirt and debris further into the coils making the situation worse. No organisation will want to publicly report their failings but one of my clients has measured a 57% improvement in energy efficiency after a full deep clean of 3 process air cooled chillers.

Dirty coils will reduce the cooling capacity, in some cases this will cause high pressure trips, which has become an industry concern with increases in peak ambient temperatures.

Refrigeration equipment is selected to achieve a designed cooling capacity based on a peak ambient temperature. System designers will select the equipment depending on the geographical region and the end user specification.

Over the last 25 years designers have had to specify systems to operate at higher peak ambient temperatures, the higher the rating the more expensive the system.

It is worth pointing out that the recorded peak ambient temperatures is measured in a defined method and varies considerably to the actual conditions. What is worrying and has not been widely reported is that some areas in the UK which have never suffered with high ambient temperatures have and that some areas have had 3
day consecutive days where the ambient temperature did not drop below 20°C.

Other things will affect condenser performance (reduce capacity) for example short cycling, low noise/speed condenser fans and poorly located equipment.

- 4 sites out of 9 sites in Group 1 and 4 sites out of 7 sites in Group 2 had issues with short cycling
- 4 sites out of 9 in Group 1 had equipment with low speed fans, one site in Group 2 had low (noise) speed fans despite being next to a major airport
- 2 sites out of 9 sites in Group 1 and 5 sites out of 7 sites in Group 2 had poorly sited condenser

Air inside the data halls is filtered before being pulled through the evaporator coils, so it was not surprising that all the coils inside the CRAC units were spotless.

It is worth remembering that dirty evaporator coils will have minimal impact on system performance whereas dirty condenser coils will have a dramatic effect on system performance.

In the last few years, the situation has been exacerbated because of the introduction of microchannel condenser coils. These coils are cheaper to make, made from aluminium and have been designed to reduce the amount of refrigerant charge required. In the early days there was an epidemic amount of coil failures, often with the entire refrigerant charge lost.

Because of the design we believe microchannel coils are more prone to fouling and require special care when cleaning.

The Institute of Refrigeration has published a good practice guide `GPG 139 Impact of Dirty Coils' which describes the correct method of cleaning condenser coils and is available free to non-members.

4 sites in Group 2 were using hybrid free cooling chillers, these units come with an additional free cooling circuit, the free cooling coil is located directly in front of the refrigeration coil. The principle of operation is when the ambient temperature is low enough no mechanical cooling is required.

In hybrid free cooling chillers, air is pulled though two heat exchanger coils, there is a small gap between the 2 coils which are fixed permanently in place. This design makes effective coil cleaning impossible and, in many cases, we have seen the free cooling coil drained and permanently removed.

In some cases, hybrid chillers have a sperate circuit to the main chilled water circuit with an additional heat exchanger which require a separate set of circulating pumps. We found the manufacturers quoted performance data is difficult to interpret and often quoted conditions that may exist for a small number of hours per annum, but the message is clear hybrid chillers offer poor performance when compared to separate dry air coolers being installed on the same chilled water circuit. We would not recommend hybrid free cooling chillers as a solution. In Group 2 we found 5 out of 7 sites had the pipework incorrectly insulated which seems an anomaly in the data centre sector.

The temperature of the hot gas leaving the compressor will be between 60°C & 120 °C and must travel to the air-cooled condenser located outside the building, any heat produced in the halls is deemed parasitic (not useful) and the pipework needs insulating.

Once the pipework leaves the building there is no requirement to insulate any of the pipework, because the.
Any insulation will derate the condensers and reduce system efficiency and should be removed (see attached GPG 125).

The pipework leaving the condenser and going back to the CRAC unit inside the building will contain liquid refrigerant which will be close to the condensing temperature and higher than the required space temperature (not useful) and so should be insulated We found 2 sites in Group 2with insulation missing on pipework inside the buildings.

Conclusions

There seemed to be a lack of understanding and care about the correct method of coil cleaning, in some cases, coil cleaning had been carried out incorrectly making the situation worse.

Most sites did not know the correct method of cleaning the condenser coils and the function was left to the service company. We have seen reports of data centres falling over and we must prepare for higher ambient temperatures; with many systems designed for a lower peak ambient temperature this makes the correct method of coil cleaning a priority.

If heat exchangers are not cleaned correctly than expected increases in the peak ambient temperature will increase energy consumption and some systems may fail on high pressure causing risk to the business.

Contact Kaizen here

Winner takes all – meeting the demands of complex, end-date critical construction programs

The demand for end-date critical infrastructure has exploded, and nowhere more so than in the data center market. New clients, evolving technologies, and unconventional business models are reshaping delivery, even as multi-billion-dollar programs are underway.

BY CHRISTIAN GOLDSMITH, SENIOR PROGRAM DIRECTOR, TECHNOLOGY AT ARCADIS

IN A SECTOR where speed to market is a commercial imperative, navigating complexity while ensuring certainty of delivery has become one of the most pressing challenges for data center developers, owners, and operators.

Adding to the pressure, today's facilities must also be built to accommodate next-generation requirements, ranging from AI training workloads and edge computing to smart manufacturing enablement and low-carbon cooling solutions. These shifting fundamentals in data center design and construction introduce mounting challenges. Exponential growth places delivery certainty at risk as requirements evolve mid-flight, new technologies are deployed without precedent, and supply chains are under pressure

Achieving delivery certainty in this environment demands a structured yet agile approach to program management, combining deep expertise with the ability to adapt quickly and confidently to ever-changing market demands.

Certainty of delivery under pressure

Data center mega-programs are defined by their speed, scale, and technical sophistication. For developers, the ultimate priority is achieving day-one readiness. Delays are not just an inconvenience, they pose significant risks to investment returns, customer confidence, and brand reputation.

To succeed, programs must be configured around securing both capability and capacity in alignment with immovable end dates. Failing to do so increases cost and risk, but could also render the entire development financially unviable.

Navigating this complexity requires agility and expertise to mobilize programs rapidly. By leveraging global experience, forming strategic partnerships, and committing to responsible delivery practices, it is possible for project teams to achieve certainty of delivery when it comes to these fast-scaling programs.

There are three main areas of risk that need to be addressed.

Capacity and capability: stretched to the limit

Unlike traditional construction programs where compromise and value engineering are common, data centers offer no room for shortcuts. These fast-moving, high-value assets require "right first time" delivery, placing pressure on both client and delivery teams.

Modern data center programs operate as fully integrated enterprises, encompassing site selection, permitting, design, procurement, build, commissioning, and day-one operations. The team is typically composed of the client organization, consultants, designers, general contractors, specialist subcontractors, and equipment suppliers.

This is further complicated by current market dynamics:

- Supply chain shortages continue to impact equipment availability, power distribution units, generators, and specialty cooling systems. Global demand is outstripping supply.
- Advanced delivery resilience strategies are becoming standard. Organizations are investing in integrated delivery models to reduce administrative burden and ensure agility.
- Immersive cooling, indirect evaporative systems, and heat reuse infrastructure are no longer future technologies - they are being deployed now to meet ESG expectations and reduce total energy footprint.

Scaling risk in a crowded market

As more data center mega-programs launch concurrently, the delivery environment grows more competitive and volatile. While known risks such as equipment delays or integration errors remain, their impact is amplified.

Programs often rely on global, virtual teams operating across geographies. Global experience and structured digital delivery environments can ensure real-time collaboration and transparency across supply chains and stakeholder groups, mitigating risks.

At the same time, cybersecurity and geopolitical concerns are prompting clients to rethink geographies and vendor selection. Cyber-resilient infrastructure and edge deployments are growing in prominence. This impacts not only security protocols but also physical build strategies.

Emerging use cases for AI and 5G are further intensifying demand for low-latency, high-bandwidth infrastructure. These workloads drive highly customized configurations and power requirements, further increasing delivery complexity.

Rethinking the enterprise model
The surge in investment into digital infrastructure, AI, and cloud computing means that data centers are competing for resources not just within their own sector, but also with industries like life sciences, energy, and advanced manufacturing.

Sustainability mandates are a central driver. With the tech sector accounting for 2–3% of global GHG emissions, stakeholders expect data centers to demonstrate real progress. This includes local water use efficiency, heat recovery, and integration with urban planning.

Traditional delivery models, where general contractors hold lump-sum risk, are under pressure. Few contractors can take on multi-billion-dollar programs unconditionally, and fewer still are willing to do so in a capacity-constrained market. The design and procurement approach must evolve: modularity, standardized design templates, and reusable components are critical to reducing lead times and increasing predictability.

A structured approach to change

In a market characterized by speed, scale, and technical complexity, traditional data center delivery models are under strain. We know that clients face growing pressure to deliver critical infrastructure against fixed timelines, while also adapting to new technologies, evolving enterprise needs, and global supply constraints.

Project delivery assurance now hinges on modular design templates, integrated delivery solutions, and increased resiliency against operational risks such as unplanned downtime. Adapting to these changing market dynamics is not optional, it's a prerequisite for success.

To help data center owners and operators evaluate their readiness to deliver fast, complex programs, the following five points need to be considered:

- **Clarity of purpose**
  Is the delivery organization truly fit for purpose? Fast-scaling data center programs require a resilient project setup that can absorb disruption, respond to innovation, and remain focused on outcome delivery. Scaling leadership, governance, and capability is critical to withstand the pressures of high-intensity delivery.
- **Client's role**
  What must the client take responsibility for? In data center mega-programs, the client cannot delegate success. Leadership, stakeholder alignment, ownership of risk, and an ability to rapidly scale internal capability all sit squarely within the client's role. These programs demand active participation and accountability.
- **Protecting the end date**
  Is the Day 1 objective truly driving program decision-making? Defining and protecting Day 1 should shape how dependencies, risks, and milestones are managed. Scenario planning, impact modeling, and clear sequencing are essential to avoiding the downstream impacts of missed dates, particularly where long-lead equipment and tenant spaces intersect.
- **Enabling effective change**
  Can the team accommodate change without losing control? A culture of readiness is more important than rigid certainty. Teams must balance assurance with flexibility, understand the real impact of change in real-time, and adopt tools and governance that support low-pain adaptation.
- **Assured decision making**
  Using data-driven tools can help project teams make smarter, faster decisions, maintaining clarity even as complexity increases. These tools support incremental improvements across design, procurement, and delivery to safeguard certainty amid changing conditions.

By applying this framework, clients can reframe their approach to delivery and identify where incremental improvements will yield the greatest impact. It helps illuminate areas where existing structures may no longer be fit for purpose, and where better coordination, clearer roles, or improved use of data can dramatically increase certainty.

Ultimately, it's about building the confidence to deliver, with the tools, insights, and leadership needed to succeed in a market where speed and precision determine competitive advantage.

Arcadis is a leading provider of data center design, engineering and consulting Services. We focus on the delivery of data centers for the world's leading data center providers, financial institutions and asset owners, including hyperscale, co-location and investor clients. Contact me to learn more.

Contact Arcadis here

Why we can't stop building data centres – but need to do it differently

Generating a digital footprint is an inevitable part of being a citizen of the developed world, and our digital trail has to reside somewhere.

BY MATT EDGLEY, CHIEF OPERATING OFFICER AT DATUM DATACENTRES

I, FOR ONE, would much prefer to know that the digital footprint I produce each day as I work online, run Teams meetings, save things to the cloud, browse the internet (and that's not even the half of it!) is being managed in a facility that is custom designed and built to be as efficient as it possibly can be. And yet according to a survey commissioned by BCS Consulting, we're facing a lack of newly constructed data centre space in the UK due to planning and power challenges.

Hot on the heels of our 16-month construction of our latest data centre facility in Manchester (MCR2), we've learned a thing or two about data centre construction. We didn't face any significant planning or power struggles named in the BCS survey, perhaps because we weren't trying to swim against the tide:

- Our data centre construction formed part of the £500m Wythenshawe regeneration project.
- We built MCR2 adjacent to our existing MCR1 facility so could capitalise on our existing connectivity network.
- The bustling Manchester tech hub has an ongoing need for new data centre facilities to support economic growth.
- We worked closely with the local council to ensure that our construction project met the needs of the region.
- We committed to putting sustainability at the forefront of the build.

Sustainable construction doesn't mean compromising on affordability

It's true that integrating the best sustainable practices into our construction didn't make the build cheaper, but it hasn't caused our costs to skyrocket either. This means that we don't have to pass any 'environmental levy' costs onto our tenants, and it has allowed us to embed sustainable design principles to create something that's resilient and sustainable, to stand the test of time.

Arguably, it didn't make the construction process any easier, either – carrying out an embodied carbon assessment at the design phase of the project to understand the environmental impact of the construction phases and life cycle of MCR2 was complex in its own right, but it was well worth it in terms of measuring and mitigating some of our environmental impact.

And alongside the sustainability initiatives that savvy clients might expect as standard, there are others that we were only able to implement at the point of construction - design features that would be of direct benefit to the wider community in which we were building. For example, an integrated heat reuse and redistribution capability is available to help warm a nearby planned development of homes - a positive consequence of our data centre being there in the first place. This is alongside important mitigation measures such as a 900m3 attenuation tank to capture and manage site run-off from heavy rainfall events, reducing the risk of localised flooding and minimising the strain on drainage exit points and downstream systems.

But there is resistance ...

The general public might be forgiven for thinking it's all doom and gloom when it comes to data centres. Most likely they've heard about data centres' considerable power or water consumption, or they'd conjure up images of expansive, ugly facilities in arid deserts using up every last drop of the precious water resources. But that's definitely only one side of the data centre coin.

A more generous (and realistic) view would be to accept and embrace the role of data centres as part of critical infrastructure - the UK's Technology Secretary officially designated data centres as Critical National Infrastructure in September 2024 (and with good reason). Perhaps even the harshest critics of data centres may be reluctant to rewind to the days when the vast majority of our daily tasks didn't leave some kind of digital footprint that has to be processed within in a data centre. So, love them or hate them - for as long as we want to store and collaborate on shared documents, book a hospital appointment, save our family photos to the cloud, and stream our favourite TV shows in the evenings, data centres are here to stay. With this in mind, we would argue that we should expend our efforts making data centres as efficient as they possibly can be instead of continuing a debate about whether data centres are the heroes or villains of the digital age.

It may not all be roses, but the alternative is worrying

So, the environmental impact of our growing need for data centres is undeniable. But what is

DCA CORPORATE PARTNER

the alternative? What would these power, water and GHG emission figures be like if we didn't have data centres; if we went back to leaving it to individuals and businesses to run their own data centre facilities, those figures would be enormous in comparison. And if we consider the extent to which things have changed in the data centre world since the huge clunky computer 'mainframes' of the 1940s, then it's realistic to expect data centres to become ever more sustainable as time goes on. Datum's Account Director, Matt McCluney, offers a simple yet powerful analogy to understand the value of off premise data centres versus individual on premise facilities:

"Think of data centres as buses – they might be large and not always aesthetically pleasing, but isn't it far better to share the ride rather than have everyone hopping into individual cars? The alternative is far less efficient, more damaging to the environment, and completely unsustainable. And as technology improves and technological advances are made, data centre operators and their clients can reap the benefits."

Moving forward

Many neighbourhoods don't want data centres, and it's easy to understand why. But we're hoping to demonstrate through our MCR2 build (and any subsequent construction projects) that it's possible to create data centres that not only have minimal negative impact on the communities in which they sit, but actually positively contribute. And whilst

data centres will never garner universal approval, our industry has a responsibility to lead by example and build trust. By working together as an industry and by contributing to the data centre community through industry bodies like the Data Centre Alliance, we can work together and learn from each other to move towards achieving the holy grail: first-rate data centres to drive economic growth with greater efficiency and reduced environmental impact.

Contact Datum here

From load to leadership: How data centres are redefining decarbonisation for industry

In recent years, data centres have come to symbolise the high energy demands that fuel our digital lives. Headlines have painted them as giant facilities with a growing carbon footprint thanks to power-hungry operations at odds with global climate goals.

BY CHARLES JOHNSTON, SUSTAINABLE ENERGY FIRST

READERS of this magazine know the reality is far more nuanced. In fact, data centres can no longer just be characterised as high-demand energy users. In late 2024, in recognition of their strategic importance, the UK government designated data centres as Critical National Infrastructure (CNI), placing them on a par with essential services like water and energy.

This reclassification brings a host of heightened operational, regulatory and security expectations. Data centre operators are under increasing pressure to balance strike a balance between more sustainable operations and operational performance and resilience. Data centres are rising to this challenge because they have (for some time) been quietly creating a blueprint for sourcing energy more sustainably without sacrificing uptime or performance. The rest of the industrial sector is finally beginning to take notice.

How are data centres improving energy efficiency?

It's true that data centres account for a significant share of electricity demand, around 2.5% in the UK (Takci et al, 2025), but this statistic is only part

dca CORPORATE PARTNER

of a bigger picture. The real story lies not in how much power is consumed, but in how it's sourced, managed, and optimised.

Over the past couple of decades, operators have dramatically improved efficiency and taken early, proactive steps to decarbonise their energy portfolios. PUE for example has reduced from an average 2.5 in 2007 to 1.56 in 2024 (Statista, 2025).

In doing so, they've began to address their own emissions, and they've created replicable models for other sectors with similar operational demands.

Data centres were among the first to use long-term renewable energy contracts, not just for cost predictability but, as a core pillar of climate strategy. Power Purchase Agreements, both physical and virtual, are now core to how many facilities manage energy risk.

What's evolved is the sophistication of those agreements. Today, operators are pushing beyond annual matching with greener energy. Companies like Google are developing 24/7 carbon-free energy models that pair every hour of data centre consumption with cleaner generation, reshaping utility expectations and energy tracking frameworks in the process.

This has major implications for other industries. Virtually every sector is now facing pressure to decarbonise Scope 2 emissions and are looking at the PPA models and energy strategies pioneered by data centre procurement teams as the mechanism to do so. Readers of this journal will already be familiar with the relentless drive for efficiency across the sector. From optimisations in cooling technology to heat reuse and ultra-low PUEs, the industry has been innovating on infrastructure since its inception against a backdrop of digital data generation demand that has grown from 33 zettabytes in 2018 to 175 zettabytes in 2025 (Reinsel et al, 2018).

The bigger picture: Data centres pose opportunity for learning

But what is increasingly clear is that these internal gains on efficiency are part of a larger ecosystem shift. The operational practices honed within data centres such as granular monitoring, dynamic load balancing, thermal control, and integrated backup, are becoming the template for decarbonising high-density energy use elsewhere.

We're now seeing similar strategies adopted in adjacent industries: hospitals improving energy resilience with on-site generation, manufacturers integrating load-shedding into process design, and commercial buildings deploying thermal storage to avoid peak pricing. Many of these innovations were first proven viable and economically rational in data centre environments.

What does net zero look like for smaller data centres?

While hyperscalers and large colocation providers have led the sector's innovations on clean energy procurement and infrastructure optimisation, smaller data centre operators face a tougher road to decarbonisation.

Limited purchasing power, tighter margins, and fewer in-house energy specialists can make it difficult for smaller operators to access the same renewable procurement mechanisms or justify the upfront investment in sustainability-driven retrofits. PPAs, for example often require long-term commitments and creditworthiness thresholds that many independent operators or regional providers struggle to meet.

To overcome these barriers, products are now entering the market which provide an alternative to PPAs and allow smaller energy users to source the same quality of renewable energy.

We are seeing the development of platforms that allow renewable energy certificates like REGOs or GoOs to now be interval-matched, typically hourly or half-hourly, allowing buyers to understand the alignment between renewable generation and demand.

Products are entering the market which are designed to democratize access to PPAs for UK businesses, like the Sustainable Energy Consortium. Data centres with electricity consumption that traditionally restricted the cost-effectiveness of PPAs can now access electricity generated by renewable assets through fixed priced contracting models.

Typically, these work through a supplier purchasing the electricity generated by a renewable asset via a PPA model themselves, before placing it into a dedicated basket which can then be accessed by their more sustainability-focused customers.

These products are born out of the same innovative way of thinking that has allowed data centres to place sustainability at the heart of their operations without operational sacrifice. That innovation hasn't always been to create something new; it's been to improve what already exists.

As digital data generation demand continues to increase with the rise of more sophisticated and ubiquitous applications for AI, there will be huge challenges for the sector in sourcing and managing energy sustainably. But, if the sector can overcome them, it is likely that the innovations developed to do so will fill a number of the blank spaces on the blueprint for overcoming the global climate challenge.

Navigating the energy and decarbonisation landscape is complex, but it doesn't have to be overwhelming. As an active member of the Data Centre Alliance, Sustainable Energy First's, expert consultants are uniquely placed to answer your hard-hitting questions, and help decarbonise your energy supply and achieve your net zero goals.

Contact SE First here


Galaxy VS
Open
Schneider Electric
Power protection
that's out of this world.
Galaxy VS UPS
Huge power in a small footprint…
The Galaxy™ VS from Schneider Electric is the best-in-class power protection that innovates against downtime. It is an efficient, modular, simple-to-deploy 3-phase UPS that delivers intelligent power protection to edge, small, and medium data centres, as well as other business-critical applications
• 99% efficient in ECOnversion™ mode
• Compact design, optimised footprint architecture
• Battery flexibility, including Lithium-ion batteries
• New patented hybrid technology
Discover the Galaxy VS UPS
www.se.com/uk
Life Is On
Schneider Electric