DCS
DATACENTRE SOLUTIONS

DEVELOPING DIGITAL INFRASTRUCTURE IN A HYBRID WORLD

ISSUE IX 2024

HTTPS://DATACENTRE.SOLUTIONS

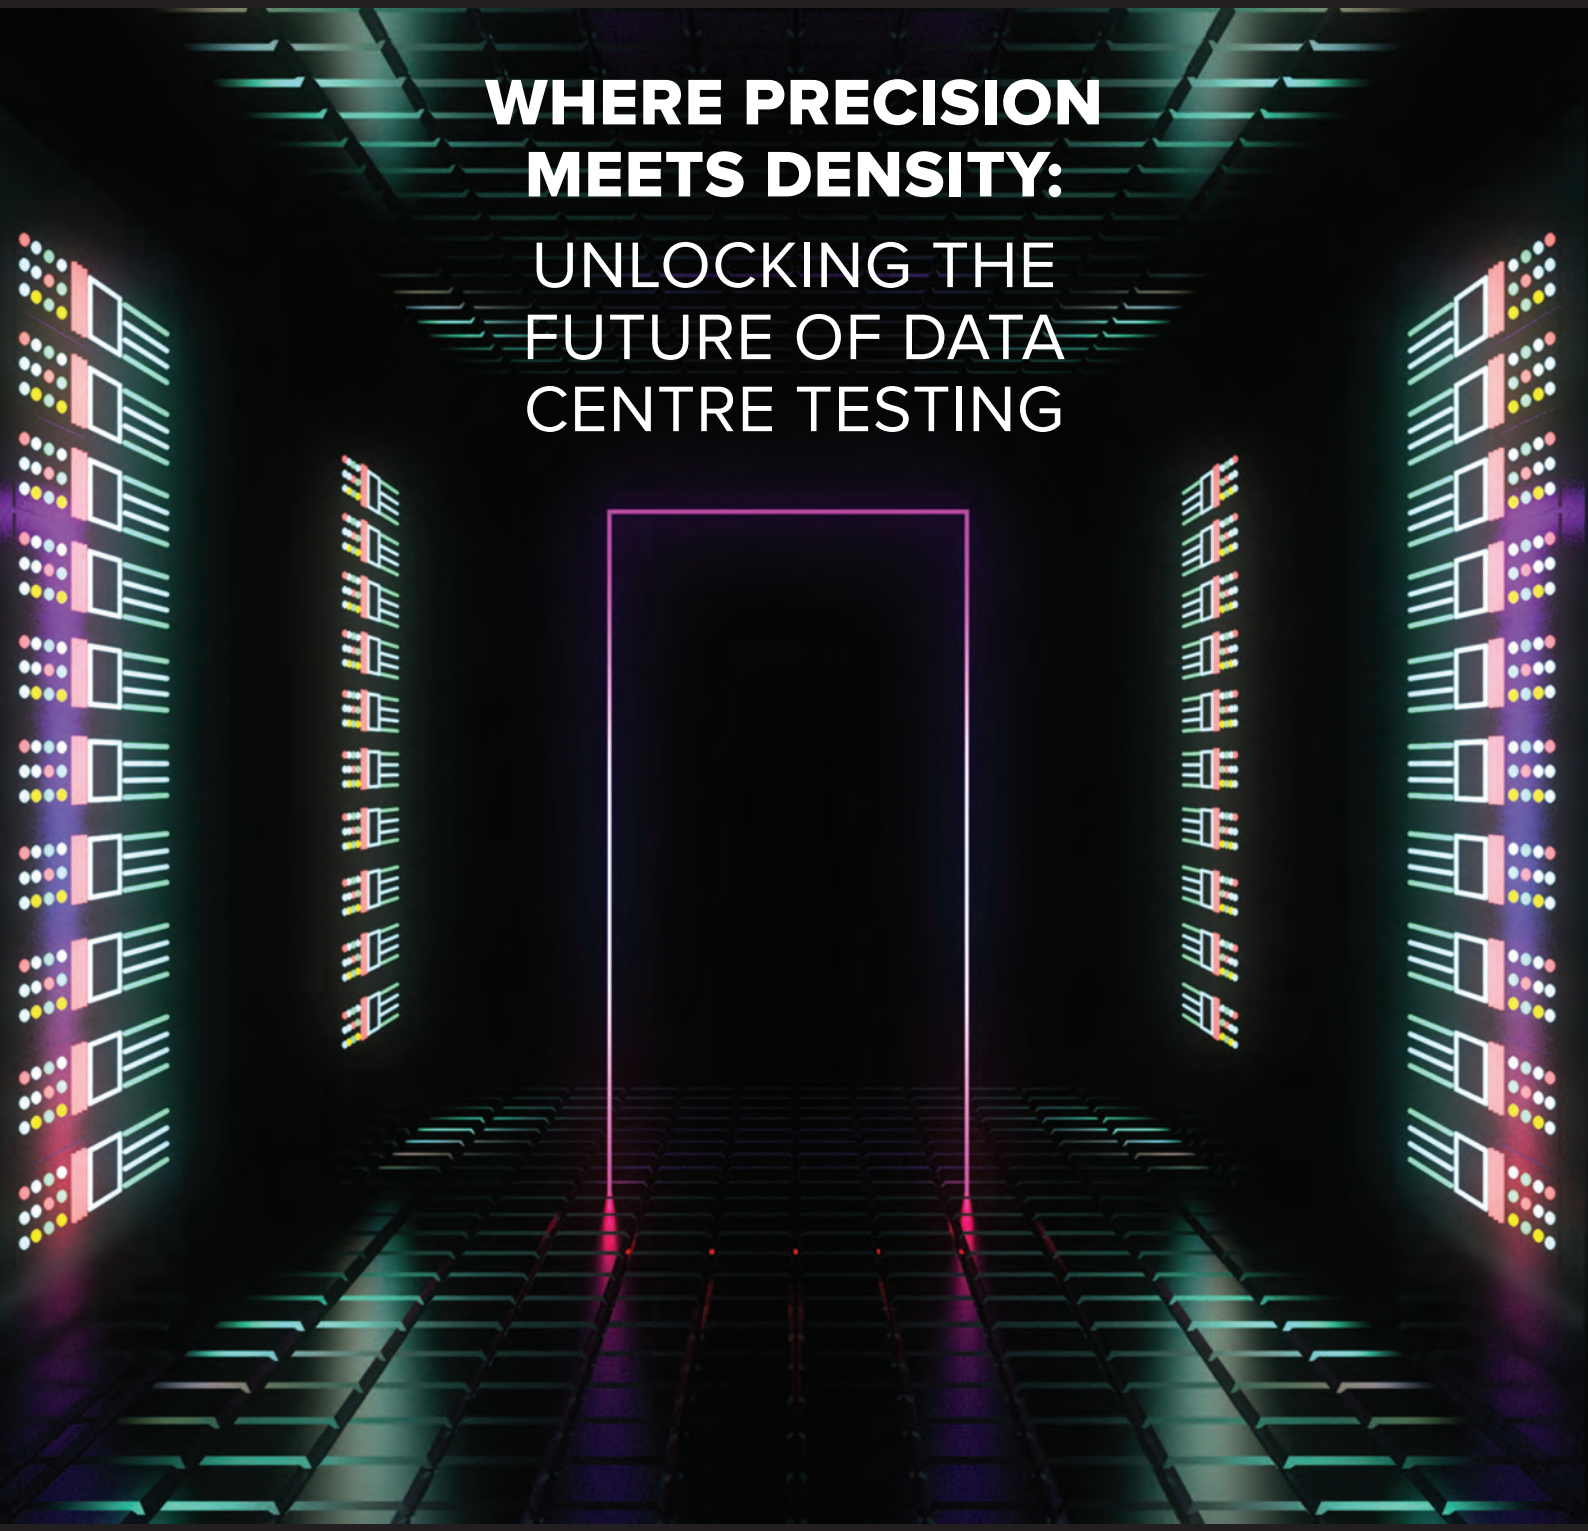

WHERE PRECISION MEETS DENSITY:

UNLOCKING THE FUTURE OF DATA CENTRE TESTING

Cloud | Design + Optimisation | Energy Management | Infrastructure + Asset Management
Hosting + Colocation | Infrastructure + Hardware | Power & Cooling

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

APC™

40Years

Make greater

IMPACT

with an industry-leading network

Celebrate 40 years with APC.

Start leveraging four decades of uninterrupted protection, connectivity, and unparalleled reliability with the APC UPS family, a legacy marked by pioneering UPS technology and an unwavering commitment to innovation.

https://www.apc.com/uk/en/product-category/88972-uninterruptible-power-supply-ups/

Life Is On Schneider Electric

VIEWPOINT

By Phil Alsop, Editor

Collaboration: crucial to future success

CHAIRING a recent roundtable on Hybrid Cooling (can be viewed via our website!), I was struck by the response to my enquiry as to whether or not the hardware folks ever actually talked to the data centre owners/operators as they developed new servers that required some form of liquid cooling (ie it might be helpful for the data centre people to know what might be heading their way to help with planning). The answer was a resounding NO. At one level, I completely understand that there is a fair distance between the design and manufacture of the latest AI-ready compute unit and an 'empty building' which is designed to house all manner of IT hardware. At another, I am somewhat surprised that there is no communication whatsoever – as new data centres are being constructed, it would be good to ensure that they are optimised to provision the necessary liquid cooling solution(s). Right now, it seems as if there's something of a guessing game going on – and the data centre owners/operators have also to be flexible enough for their customers to change their minds mid-build...

Bearing in mind that optical computing and quantum computing are likely to have an impact at some stage in the future, beyond the current AI-fuelled high-density compute 'craze', it would seem eminently sensible for the data centre industry to engage at some level with the IT hardware industry, and the server companies in particular. If only so that both parties could understand the challenges and opportunities which they provide to one another, in the hope that there might be a better way of doing things and, especially, doing some roadmap planning together.

Right now, I suspect that the Open Compute Project (OCP) is a close as we get to such cross-discipline cooperation. Let us hope that this continues and expands over time.

The roundtable discussion also touched on the importance of being able to harness the waste or, more accurately, rejected heat from the data centre – liquid cooled or not. And there was general agreement that,

moving forward, the industry had a major opportunity to build data centres in locations where they could achieve some kind of hero status. Providing district heating from this rejected heat, and also providing power to the local community via a microgrid (which could be linked up with a national power grid as well). All of a sudden, data centres have been transformed from 'power hungry polluters' to sustainability champions...

Okay, so life is never that simple. And there's quite a leap of faith to believe that governments at whatever level will understand this win/win concept. Not to mention the NIMBYs who would have us all living in some unspecified time in the past, when 'things were so much better' – unaware of the irony that much of their campaigning takes place on digital platforms. Nevertheless, we can dream, if not hope, that some joined up thinking could deliver community data centre hubs across a country. And we have the example of a company such as Deep Green to see that it can work.

All that remains is for me to wish everyone a great festive season and to look forward to a new year where, just maybe, the data centre industry gets on the front and celebrates its sustainability and societal successes, and looks to go further and faster on both into the future.

CONTENTS

ISSUE IX 2024

DCS DATACENTRE SOLUTIONS

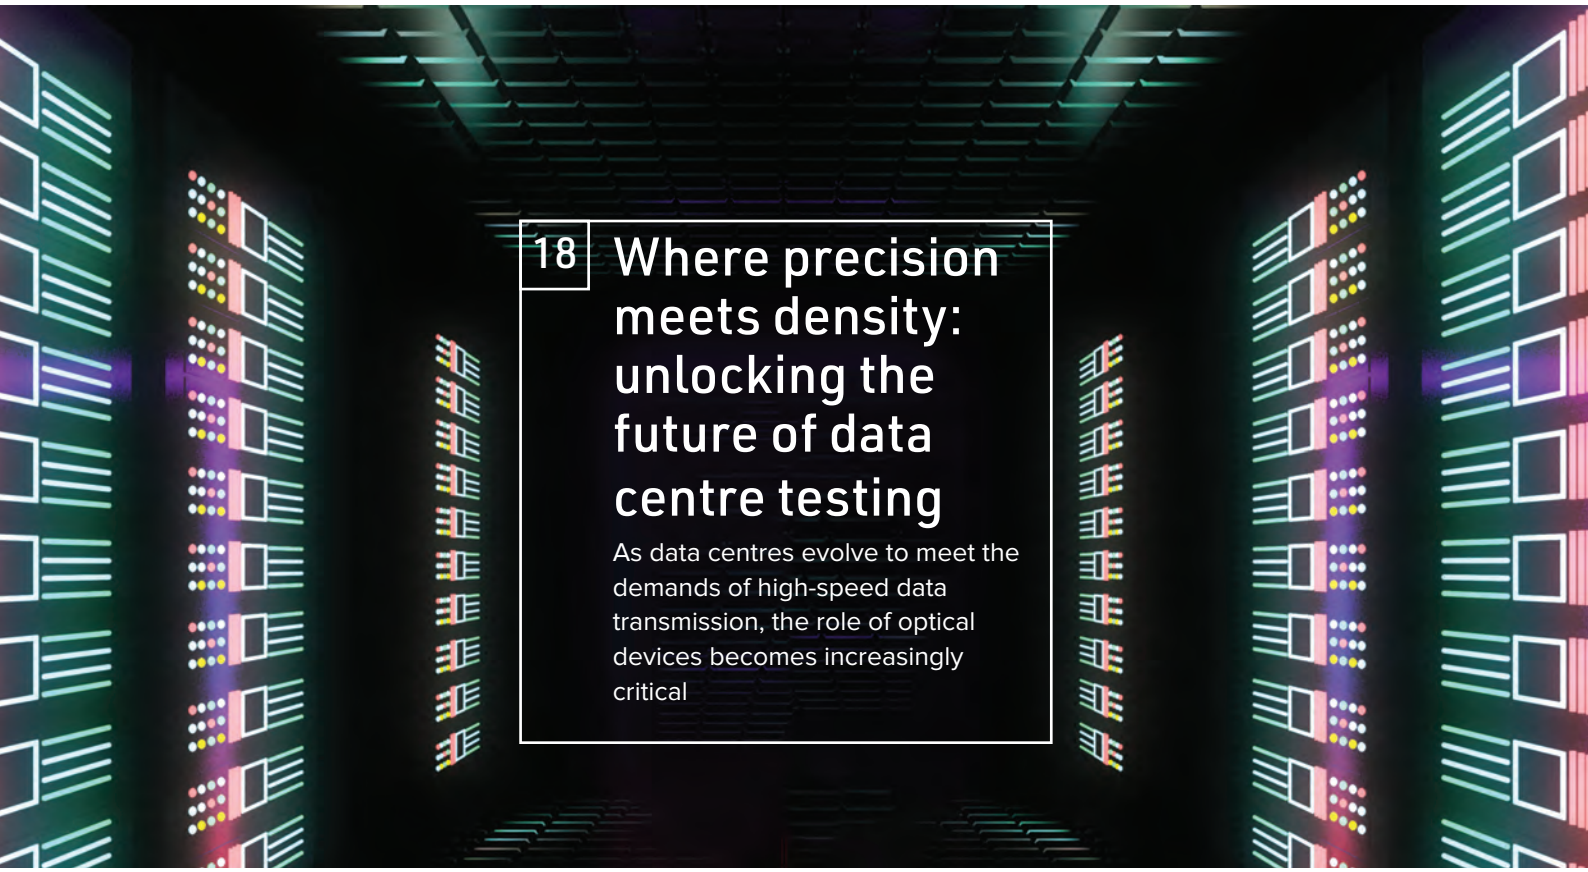

The data centre trade association

DCA News

NEWS

DCS DATACENTRE SOLUTIONS

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

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

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

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

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

Marketing & Logistics Executive
Eve O'Sullivan
+44 (0)2476 823 123
eve.osullivan@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

Datacentre Solutions is published 6 times a year on a controlled circulation basis in Europe, Middle East and Africa only. Subscription rates on request. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd. ©

ISSN 2756-1143 (Online)

The Paris Agreement's goals are no longer achievable

Despite impressive strides in 2023 and positive projections for 2024, the pace of renewable development isn't fast enough.

CAPGEMINI has published the 26th edition of its annual World Energy Markets Observatory (WEMO), created in partnership with Hogan Lovells, Vaasa ETT and Enerdata. The report takes stock of the current state of the energy transition. Despite progress being made, greenhouse gas (GHG) emissions are continuing to increase, reaching a new record high of 37.4 billion tonnes (Gt) in 2023[1], confirming that the path to the reach Paris Agreement's objectives is not on track. The report provides insights on what the key focus areas would need to be, moving forward, to address the complex energy transition challenges, including a change in the measurement of clean energy progress, as well as accelerated investment in the power grid and clean technologies.

James Forrest, Global Energy Transition & Utilities Industry Leader at Capgemini says: "Despite an historical spike in renewable penetration, the pace of development isn't fast enough to close the gap. There is still much to do in the next decade to get closer to net zero by 2050 and achieve a successful energy transition: whether it be in the field of low carbon technologies, R&D efforts, nuclear or grid flexibility and storage.

In addition, beyond the necessary adoption of new market mechanisms, a shift away from measuring energy based on primary consumption is needed. This measurement was relevant during past energy crises, but it is now time to adopt a more holistic approach. Moving to a final energy demand measurement would better assess clean energy progress and ensure more accurate projections."

There is a need to hasten the deployment of renewable energy globally, and to accelerate in

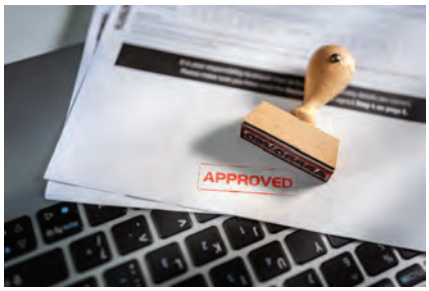

developing countries, to deliver the 2030 and 2050 decarbonization goals. The total amount of final energy provided by renewable energy is likely to be limited to about 40% of global needs. In 2023, total renewable energy capacity increased by 14% year on year with a larger capacity expansion of solar (32%) than wind (13%).

But, whilst 2024 is promising to hit another record, as this was the case for the 22nd previous years, this growth is far below what is needed to achieve net zero carbon in 2050. Moreover, while the renewable penetration rate increases, they are impacting grid stability and association with stationary batteries will become compulsory. According to the report, storable renewable energies development, such as biomass or geothermal energy, should be accelerated.

Hydrogen is now a strategic lever in the decarbonization path. The number of projects reaching final investment decision has quadrupled over the last two years. However, a refocus of applications has been observed due to the increasing costs of low-carbon hydrogen production, competition between uses, and regulations.

Only certain uses in 'Hard to Abate' industries, such as heavy industry and maritime mobility, have strong potential. Global nuclear capacity needs to triple to ensure stable, low-carbon power.

COP28 has recognized the critical role of nuclear energy for reducing the effects of climate change. While there is some promising progress in nuclear renaissance, including Small Modular Reactors (SMRs), development of new nuclear power plants is still difficult. In 2023, 440 nuclear reactors (390 GW) provided 9% of the world's electricity, 25% of the world's low-carbon electricity.

SMRs are in the planning or early construction stages with many years before they are deployed at scale as their industrialization can prove to be complex. According to the report, more focus needs to be placed on extending the life of existing nuclear plants. The power grid plays a fundamental role to accelerate clean energy transitions.

Grid investment is starting to pick up and is expected to reach USD 400 billion in 2024[2], with Europe, the United Sates, China and parts of Latin America leading the way. According to the report, better forecasting electricity consumption and finer optimization scenarios thanks to technologies such as AI will help to improve grid balancing.

Whilst AI has the potential to significantly accelerate decarbonization, a lack of skills and a focus on short-term proof of concepts is hampering adoption to date. However, AI coupled with GenAI in agentic LLM (Large Language Model) workflows[3] has a clear role to play as a catalyst to improve grids efficiency, e-fuel discovery; new battery or wind turbine design; synthetic biology; and augmented insights from many data sources for better informed decision making.

Energy efficient data centres could save up to €25 billion by 2030

Report highlights impact of hyperconverged infrastructure, hybrid cloud and leveraging on-demand computing capacity.

NUTANIX has published the findings of a new report focused on improving sustainability in data centres. Atlantic Ventures' report, Improving Sustainability in Data Centers 2024, reveals how next-generation data centre architectures, including hybrid cloud and hyperconverged infrastructure (HCI), can significantly reduce energy consumption, lower carbon emissions, and drive cost savings across the EMEA region. In just six years, the report finds that modernising data centres with HCI-based solutions could save up to 19 million tCO2e in the EMEA region, equivalent to the emissions of almost 4.1 million cars. It could also save €25 billion by 2030 from improved energy and operational efficiencies.

As businesses face-up to post-Covid digitisation and the demands for data-hungry technologies, such as AI and IoT, the report identifies an increasing urgency for action. The dual challenges of rising energy costs and increased regulatory pressure to reduce their environmental impact are making this more difficult for IT leaders. As a result, energy efficiency has become a top priority for CIOs and data centre managers.

"Data centres are critical to the global digital economy but also rank among the largest consumers of energy," says Sammy Zoghlami, SVP EMEA at Nutanix. "In EMEA alone, data centres demand over 98 TWh of energy annually, equivalent to the consumption of an entire country like Belgium. The findings of this report show that by leveraging HCI-based solutions, companies can make a powerful contribution to climate action while significantly cutting operational costs." "CIOs and digital executives are facing challenges to provide the digital infrastructure to cope with the fast-growing demand for compute power and storage capacity, especially with the emergence of AI applications," says Carlo Velten from Atlantic Ventures. "As IT budgets are under pressure and electricity prices are soaring, energy efficient data centre and cloud operations are key levers for profitability and sustainability. Hyperconverged infrastructure is at the forefront of transforming data centres into more energy-efficient and climate-friendly operations, as this report confirms."

Key findings from the report including UK specific data:

- **27% Energy Savings -** Switching from traditional 3-Tier architectures to an HCI-based platform can reduce energy consumption by more than 27% annually, helping companies cut both operational costs and emissions.
- **Massive Regional Impact -** Across the EMEA region, a full-scale transition could save up to 92 TWh of electricity and eliminate 19 million tons of CO_2e between 2024 and 2030 - comparable to the emissions of 4.1 million cars. In the UK alone, this amount would equate to savings of 13.4 TWh of electricity.
- **€25 Billion in Savings -** The financial windfall from reduced electricity consumption could reach €25 billion by 2030, offering businesses a rare opportunity to align sustainability with profitability. Due to relatively high price levels for electricity the potential electricity cost savings for companies and service providers in the UK could be as much as 3,3 billion euros between 2024 - 2030 when switching on premise from 3-Tier to HCI.
- **HCI in the Cloud -** Migrating HCI platforms to co-location or public cloud environments magnifies these benefits, with potential energy savings reaching as high as 54% compared to traditional on-premise data centres. This is due to the low PUE of Public Cloud Providers as well as the flexibility in providing on-demand computing capacity.
- **Disaster Recovery Efficiency -** HCI-based architecture also enables lean, energy-efficient disaster recovery systems in the Cloud, reducing the infrastructure footprint while maintaining scalability and responsiveness.

Dublin data centre denial highlights growing challenges

The recent denial of a Google Data Center at Grange Castle Business Park in Dublin highlights the growing challenges in powering and constructing new data facilities.

AS POWER HUNGRY technologies such as artificial intelligence (AI) expand, there is an increasing need for sustainable power solutions and enhanced data connectivity. Balancing the demands of AI with environmental and infrastructure constraints is crucial for governments and service providers, says GlobalData, a leading data and analytics company.

Gary Barton, Research Director for Enterprise Technology and Services at GlobalData, comments: "Creating favorable conditions for major tech companies to locate their European HQs and mission-critical facilities in and around Dublin has been a major tenet of Ireland's recent economic strategy so the importance of this decision should not be underestimated. However, the huge power demand of data centers is also a critical problem for national power networks around the world, creating a further headache for often outdated power grids that are also struggling to adapt to other modern phenomena such as electric cars."
The carbon impact of AI is also another major factor at play, both in the decision made in Ireland and for the industry globally. Carbon footprint deadlines are looming, with many set for 2030, and yet the need to stay competitive and find new revenues streams is pushing enterprises towards AI.

Companies such as Google, Microsoft, and Amazon are keen to highlight their heavy use of renewable power at their facilities. However, global supply of green power is limited and lagging demand, and there are questions about greenwashing when it comes to powering data centers.

Barton continues: "Many sites that claim to use 100% renewable energy are actually reliant on carbon offsetting

to achieve their targets. While not inherently incorrect, this practice is facing growing scrutiny from those evaluating carbon footprint claims. Consequently, data center operators are increasingly considering locations near carbon efficient power generation facilities such as nuclear power plants. The decision in Dublin is likely to be the first of many that will ultimately require new data center builds to directly or indirectly fulfil more of their own power requirements either on-site or nearby – and this will have to be low-carbon power."

Power is not the only requirement for data centers where supply is outstripping demand. AI solutions have turbocharged the growth of data going into and out of data centers and the connectivity required to support this growth is not automatically available. Barton concludes: "Network building service providers and subsea cable consortia are unsure where AI facilities are likely to be located and this adds uncertainty to the future network build out planning. While hyperscalers and service providers have made agreements to secure future data capacity, there is no guarantee these arrangements will fully meet demand. AI holds immense potential, yet it is essential for governments and providers to collaborate to ensure sufficient capacity while minimizing impacts on national infrastructure and the environment."

Rising demand for data centres prompts innovative cooling solutions
The rapid growth of artificial intelligence (AI) has significantly increased the demand for data processing capabilities and capacity, leading to the expansion of data centers globally. As these facilities operate continuously, their energy-intensive equipment are raising concerns about sustainability. New cooling technologies and innovative approaches are emerging to address the environmental challenges posed by these critical infrastructures, says GlobalData, a leading data and analytics company.

GlobalData's latest Strategic Intelligence report, "Deep Dive into The Environmental Impact of Data Centers," reveals that water is the most common cooling method, and its use has increased during the current AI boom. However, new cooling technologies are being developed.

Available for sales and rental

CRESTCHIC

THE **TRUSTED** NAME FOR **LOAD TESTING**

200kW Data Centre Loadbank: Built for Integrated Systems Testing

Simulate full load operation, power outages, and cooling system failures

Verify PDU operational readiness during commissioning

Closely mimic server heat generation to test cooling systems

Control up to 20 units from one controller

- Small footprint, stackable for easy storage
- AC load bank resistive-only (200kW)
- Voltage range: 400V or 415V
- 50/60 Hz Operating frequency range
- 3 Phase
- Forklift pockets for lifting
- Sheet steel construction with IP55 rating
- Low Delta T
- Low noise
- Internal or external auxiliary supply operation

www.crestchicloadbanks.com

sales.uk@crestchic.com | hires.uk@crestchic.com | service.uk@crestchic.com

Sales & Rental Offices in : UK, North America, Singapore, Germany, Belgium, France, China and UAE

Majority of data centre businesses confident in their energy strategies

A significant 85 percent of UK data centre businesses report confidence in the cost-effectiveness, reliability, and environmental sustainability of their energy strategies, according to research conducted by True, powered by Open Energy Market. The findings, detailed in a new whitepaper, 'Mind The Gap: From Ambition To Action', explore the commercial imperatives driving energy procurement and sustainability strategies among UK businesses.

WHILE MANY ORGANISATIONS report high confidence in their current strategies, underlying challenges threaten to undermine these efforts, exposing a risky "set and forget" mentality that could have significant implications for hindering commercial benefits as well as the future of the UK's sustainability efforts.

Almost all UK data centre businesses (89 percent) believe they are prepared to tackle challenges such as energy market volatility and tighter environmental regulations within the next 12 months.

Despite this reported confidence, there is a concern that businesses may be overestimating their readiness, leading to complacency and missed opportunities for optimisation. For example, this overconfidence is contrasted by the 31 percent of data centre businesses reporting integration of new data-led technologies within existing systems is a significant hurdle. Brand reputation and public image has also been identified as a key concern by four in ten (41 percent) data centre businesses, much higher than the overall national average (27 percent) for UK businesses.

Chris Maclean, CEO, True, powered by Open Energy Market, comments: "This shouldn't be about competing priorities; it's about recognising that sustainability can drive commercial benefits. As part of a high-visibility, energy-intensive industry, data centre businesses need to shift their perspective to see sustainability as a pathway to commercial gain, or they risk losing

competitive advantage, missing out on big tech or corporate tenders, and falling behind in an evolving market landscape where sustainability is increasingly tied to business success". Plus, data centre businesses report common challenges for implementing sustainable energy practices. These include limited availability of renewable energy sources (29 percent), financial constraints (25.5 percent), followed by extended timelines for approval and decision-making processes and operating in silos with fragmented net zero plans (23.5 percent). These obstacles underscore the need for streamlined processes and better access to expertise to drive effective sustainability initiatives that enhance commercial performance.

The introduction of new technology to monitor and improve energy usage is a prominent strategy, adopted by 43 percent of data centre organisations. However, from a strategic standpoint, the survey reveals a complex landscape of decision-making authority over sustainable energy strategies. The data centre industry is in a unique position compared to the rest of the sectors surveyed (Food and Drink, Manufacturing, and Hospitality), where Heads of Sustainability edge out CFOs as the final sign-off authority (41% percent versus 13 percent).

Still, almost half of data centre organisations (45 percent) report a 4-6-month timeline for signing off energy and sustainability projects, with 23 percent taking 7-9 months. Lengthy approval processes may hinder timely implementation of crucial initiatives.

Moreover, the research reveals limited evaluation methods for sustainable energy transition action plans. Only 17 percent of data centre organisations track and analyse energy consumption data, and similarly few review compliance with environmental regulations (17.6 percent) or compare performance against industry standards (15.6 percent).

True provides a comprehensive solution to these challenges. By unifying energy procurement and net zero strategies in one platform, True empowers businesses to make data-driven decisions and build stronger financial grade business cases for their sustainability initiatives.

80% of US Internet Exchanges are now data centre and carrier neutral

Independent study reveals US trend toward data centre and carrier neutrality to support new AI, cloud, and IoT applications.

THE US is steadily moving toward a data center and carrier neutral framework to support new AI and cloud-based applications, with 80% of all Internet Exchanges (IXs) in the region now data center and carrier neutral. That is according to a new independent study conducted by Dstream Group on behalf of DE-CIX, the global Internet Exchange operator.

With record-levels of net absorption of data center capacity throughout the US, the impact of the generative AI boom is being felt across the board. This comprehensive study provides insights into how the US interconnection landscape has evolved to cope with this demand and provides in-depth predictions about its future. In response to the rising need for high-performance interconnection – driven by the rise of cloud computing, artificial intelligence, and interest in online gaming and high-resolution streaming – deployment of IXs in the country has surged by 600% in the past decade, the majority of which are now data center and carrier neutral.

These neutral platforms have on average four times more data centers from various operators connected to the platform than other IX models, offering enterprises and network operators greater choice and resilience opportunities.

Operated by independent specialists and distributed over data centers from multiple operators within a city or region, neutral IXs offer more access points and bring together more networks than the previously prevailing data center (DC)/carrier operated IXs. Of the top 50 largest IXs in the US, 35 (70%) are neutral, demonstrating a strong preference of network operators for the model. Neutral IXs create

network density and opportunities for businesses and operators to leverage more resilient, low-latency, edge-based connectivity.

The study also reveals that in today's top 50 US-based IXs, the distributed and neutrally operated IXs have an average of 11 connected facilities operated by a minimum of two independent data center operators within a metro area. In comparison, the DC/Carrier-operated IXs have an average of three facilities operated by a single operator within a metro area. Greater operator diversity and geographical distribution within the metro area allows companies freedom of choice of data center operators and brings interconnection closer to end-users, reducing latency and improving connectivity performance.

"The distributed, independently operated IX model has several significant advantages for building digital ecosystems," explains Serge Radovcic from Dstream Group, co-author of the study. "The neutral model can potentially be accessible from all colocation data centers within a metro area – and even from outside of the metro area. By leveraging connectivity to multiple data center operators, an IX can eliminate the risk of vendor lock-in, and make it easier to establish redundant connections, increasing the resilience of connectivity for critical use cases."

How neutral IXs encourage ecosystem growth is exemplified in the study on the basis of New York, an interconnection powerhouse. New York is unique in the US interconnection landscape for having two large-scale neutral IXs alongside two large DC/carrier operated IXs, and almost double the average number of networks per IX found in US cities (134 compared to 70).

In New York, this density is also reflected in the surge in demand for colocation space. The distributed and neutral IX model enables data centers outside of the city center to be incorporated into the interconnection ecosystem, providing companies with solutions to the data center squeeze being experienced in some historical hubs.

The study also investigates the growth in the data center market in the US, demonstrating the link between strong neutral IXs in a market and a diverse ecosystem of connectivity-focused data center operators.

Demand is high and vacancy rates are low. Reflecting this, data center planning and construction is increasing, with Northern Virginia and Dallas/Fort Worth in the lead. With a current total of 11,200 MW of data center installed capacity across the US, a further 5,500 MW is currently being built, and another 12,600 is in planning, coming to over 160% growth in the near future.

As future trends in the interconnection and data center markets noted in the study – such as edge computing, disaggregated computing, and AI – continue to evolve, the demand for the resilience and flexibility offered by neutral IXs will only grow.

Hybrid hosting now preferred choice for IT leaders

Research released recently shows that 67% of IT decision makers favour a hybrid hosting infrastructure over a “cloud-first” strategy and 94% of businesses using public cloud would take a different approach to migration, given the challenges they have faced when moving to the public cloud.

THE REPORT, commissioned by Asanti, one of the UK’s leading colocation data centre providers, and conducted by VansonBourne, indicates that many organisations now view blending public cloud with on-premise solutions as the optimal approach to managing IT infrastructure.

Over the last decade, many organisations adopted a cloud-first approach, however as the research findings reveal the realisation that a hybrid strategy, combining on-premise and cloud, would have helped to avoid some of the critical challenges shared in the report. The research, which surveyed 100 senior IT decision makers across public and private sectors in the UK, revealed key drivers for this change in thinking:

- **Unexpected costs:** 77% of respondents reported that operating costs in the public cloud were higher than anticipated, with 63% stating these costs exceeded those of their previous non-public cloud models.
- **Security and compliance concerns:** Ensuring security and meeting regulatory compliance proved challenging for 62% of organisations. As a result, many are reconsidering the public cloud for their critical and sensitive data.
- **Prolonged migration times:** 98% said that they faced difficulties during their migration to the public cloud with more than half (57%) reporting that the move to public cloud took longer than expected, often leading to disruptions in business operations and escalating costs.

These challenges have led to widespread repatriation of applications from public cloud platforms, with a staggering 91% of those surveyed reporting that they are now moving applications back to on-premise or

colocation data centres. Performance issues, downtime, and the need for more control over infrastructure are additional reasons for this shift. An over-reliance on support from public cloud providers, cited by 60% of respondents as their number one resource for migration planning, may played contributed to these issues.

Other significant findings related to public cloud migrations include:

- 47% experienced service downtime and business interruptions during their public cloud migrations.
- 44% indicated a need to upskill internal teams to manage and optimise public cloud infrastructure.
- 38% noted performance degradation after migrating to the public cloud, with many finding cloud resources insufficient to meet their performance expectations.
- 31% faced application compatibility issues, complicating their IT integration efforts.

Stewart Laing, CEO of Asanti, commented: “With such a high percentage of organisations bringing applications back to on-premise or colocation data centres, you have to wonder – what exactly were cloud providers promising when they sold these packages? It’s clear now that defaulting to a cloud-first strategy may not be the best approach. We strongly encourage IT decision makers and business leaders to adopt a business-centric view, focusing on current and future needs to determine the most appropriate hosting model, rather than starting with a predefined solution and attempting to fit the business around it.”

The report also highlighted a growing focus on AI applications with 52% of organisations stating that the adoption of AI within their business would require a change to their public cloud strategy, with on-premise and/or in-house IT infrastructure (colocation)) preferred to public cloud. This approach will enable businesses to feel assured that they are operating within a stable and secure environment for sensitive workloads.

AFL

“The emergence of generative AI, with its exceptionally large models and truly extraordinary computing requirements, has driven greater and more rapid change in data center networking in the past two years than I have seen in the previous decade.”

Alan Keizer
Senior Technology Advisor, AFL

www.aflglobal.com

AI drives data centre energy consumption increase

Datacenters are the heart of the digital economy and the demand for datacenters is expected to rise substantially, positioning them as a primary focus for growth and investment. But the cost of operating datacenters is also rising substantially due to rising electricity prices and increased datacenter consumption. A new International Data Corporation (IDC) report looks at datacenter electricity spending and the implications for technology providers and datacenter operators.

ELECTRICITY is by far the largest ongoing expense for datacenter operators, accounting for 46% of total spending for enterprise datacenters and 60% for service provider datacenters. And electricity consumption is growing rapidly as datacenters take on more workloads and more energy-intensive workloads, such as artificial intelligence.

IDC expects the surging demand for AI workloads will lead to a significant increase in datacenter capacity, energy consumption, and carbon emissions, with AI datacenter capacity projected to have a compound annual growth rate (CAGR) of 40.5% through 2027. Accordingly, AI datacenter energy consumption is forecast to grow at a CAGR of 44.7%, reaching 146.2 Terawatt hours (TWh) by 2027 with AI workloads consuming a growing portion of total datacenter electricity use.

Overall, IDC expects global datacenter electricity consumption to more than double between 2023 and 2028 with a five-year CAGR of 19.5% and reaching 857 Terawatt hours (TWh) in 2028. At the same time, electricity prices are rising due to supply and demand dynamics, environmental regulations, geopolitical events, and sensitivity to extreme weather events fueled in part by climate change. IDC believes the trends that have caused electricity prices to increase over the last five years are likely to continue. Rising consumption and increased energy costs will make datacenters considerably more expensive to operate, but how much is uncertain.

To better understand the impact of rising electricity costs on datacenter operations, IDC conducted scenario planning for a datacenter with 1 MW of IT load in 2023, running at 50% capacity and power usage effectiveness (PUE) of 1.5. The study looked at three energy price growth scenarios using energy

pricing and growth rates for the United States, Germany, and Japan. In all three scenarios, the percentage growth in electricity spend exceeds a CAGR of 15% in all cases, with most scenarios showing growth of over 20%. The study also shows that an additional 10% in energy efficiency can offer considerable savings to datacenter operators.

"There are any number of options to increase datacenter efficiency, ranging from technological solutions like improved chip efficiency and liquid cooling to rethinking datacenter design and power distribution methods," said Sean Graham, research director, Cloud to Edge Datacenter Trends at IDC. "But providing energy-efficient solutions is only part of the equation for meeting customer needs. Datacenter providers, including cloud and colocation services, should continue to prioritize investment in renewable energy sources. By investing in renewables, they are helping to increase overall supply while helping their customers meet their sustainability goals."

Solar and wind power, in particular, offer significant environmental advantages while also providing the lowest levelized cost of electricity (LCOE), which reflects the average net present cost of electricity generation over a generator's lifetime. And by collocating facilities at or near the source of renewable energy generation, providers can reduce both construction costs and energy losses associated with distribution, enhancing overall efficiency and sustainability while also improving resiliency by removing grid reliability issues.

Shared Cloud Infrastructure continues to lead infrastructure spending

According to the International Data Corporation (IDC) Worldwide Quarterly Enterprise Infrastructure Tracker: Buyer and Cloud Deployment, spending on compute and storage infrastructure products for cloud deployments, including dedicated and shared IT environments, increased 61.5% year over year in the second quarter of 2024 (2Q24) to $42.9 billion. Spending on cloud infrastructure continues to outgrow the non-cloud segment with the latter growing by 41.4% in 2Q24 to $19.4 billion. The cloud infrastructure segment experienced lower growth in unit demand at 17.7%, due to a continued increase in average selling prices (ASPs), mostly related to the exponential increase of GPU server shipments.

"Cloud infrastructure spending growth continues to be driven by accelerated AI-related investments, which especially impacted servers but also triggered enterprise storage spending," said Juan Pablo Seminara, research director, Worldwide Enterprise Infrastructure Trackers. "Different surveys conducted by IDC in 2024 show how AI investment plans have been scaling up and driving investment priorities for almost every region. Hyperscalers, Digital Service Providers, and major Cloud Service Providers are the ones that keep pushing the growth and that will continue to have a positive impact on the market during 2024 and 2025. And the improved economic prospects will help to extend the positive mood even further."

Spending on shared cloud infrastructure reached $35.3 billion in the quarter, increasing 74.9% compared to a year ago. The shared cloud infrastructure category continues capturing the largest share of spending compared to dedicated deployments and non-cloud spending, with shared cloud accounting for 56.6% of the total infrastructure spending in 2Q24. The dedicated cloud infrastructure segment presented lower growth of 19.2% year over year in 2Q24 to $7.6 billion.

For 2024, IDC is forecasting cloud infrastructure spending to grow 48.8% compared to 2023 to $164.0 billion. Non-cloud infrastructure is expected to grow 11.7% to $67.5 billion. Shared cloud infrastructure is expected to grow 57.9% year over year to $131.9 billion for the full year. Spending on dedicated cloud infrastructure is also expected to have double-digit growth in 2024 at 20.4% reaching $32.1 billion for the full year. The subdued growth forecast for non-cloud infrastructure at 11.7% in 2024 reflects that even though most of the growth will come from cloud spending, general non-cloud dedicated systems are consolidating their recovery this year.

IDC's service provider category includes cloud service providers, digital service providers, communications service providers, hyperscalers, and managed service providers. In 2Q24, service providers as a group spent $41.8 billion on compute and storage infrastructure, up 64.2% from the prior year. This spending accounted for 67.2% of the total market. Non-service providers (e.g., enterprises, government, etc.) also increased their spending to $20.5 billion growing 38.2% year over year. IDC expects compute and storage spending by service providers to reach $157.8 billion in 2024, growing at 49.4% year over year.

On a geographic basis, year-over-year spending on cloud infrastructure in 2Q24 showed very positive

results across all regions where the fastest growing regions were Asia/Pacific (excluding Japan and China), Japan, USA, and Canada with 110.7%, 98.1%, 72.1% and 53.8% year-over-year growth respectively. The other regions also showed very solid yearly increases with Central & Eastern Europe, Western Europe, China, Middle East & Africa, and Latin America growing by 48.7%, 27,7%, 24.8%, 23.4% and 9.7% in that order.

Long term, IDC predicts spending on cloud infrastructure to have a compound annual growth rate (CAGR) of 18.1% over the 2023-2028 forecast period, reaching $253.0 billion in 2028 and accounting for 76.4% of total compute and storage infrastructure spend.

Shared cloud infrastructure spending will account for 78.6% of the total cloud spending in 2028, growing at a 18.9% CAGR and reaching $198.8 billion. Spending on dedicated cloud infrastructure will grow at a CAGR of 15.3% to $54.3 billion. Spending on non-cloud infrastructure will also rebound with a 5.3% CAGR, reaching $78.3 billion in 2028. Spending by service providers on compute and storage infrastructure is expected to grow at
a 17.1% CAGR, reaching $233.0 billion in 2028.

AI requires engineering workforce to upskill

Through 2027, generative AI (GenAI) will spawn new roles in software engineering and operations, requiring 80% of the engineering workforce to upskill, according to Gartner, Inc.
"Bold claims on the ability of AI have led to speculation that AI could reduce demand for human engineers or even supplant them entirely," said Philip Walsh, Sr Principal Analyst at Gartner.
"While AI will transform the future role of software engineers, human expertise and creativity will always be essential to delivering complex, innovative software."

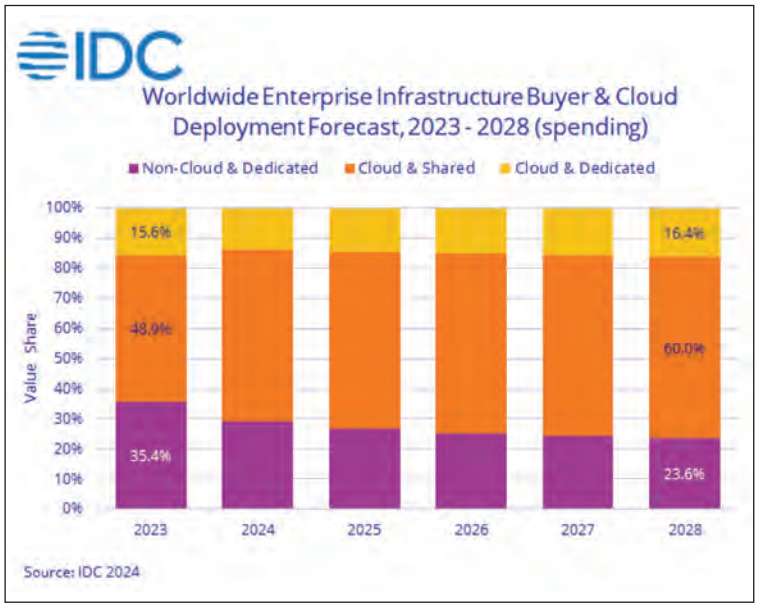

A graph of data on a white background Description automatically generated with medium confidence

Gartner analysts expect AI will impact the software engineering role in three ways:

In the short term, AI will operate within boundaries

- AI tools will generate modest productivity increases by augmenting existing developer work patterns and tasks. The productivity benefits of AI will be most significant for senior developers in organizations with mature engineering practices.

In the medium term, the emergence of AI agents will push boundaries

- AI agents will transform developer work patterns by enabling developers to fully automate and offload more tasks. This will mark the emergence of AI-native software engineering when most code will be AI-generated rather than human-authored.

"In the AI-native era, software engineers will adopt an 'AI-first' mindset, where they primarily focus on steering AI agents toward the most relevant context and constraints for a given task," said Walsh. This will make natural-language prompt engineering and retrieval-augmented generation (RAG) skills essential for software engineers.

In the long term, advances in AI will break boundaries and will mark the rise of AI engineering

- While AI will make engineering more efficient, organizations will need even more skilled software engineers to meet the rapidly increasing demand for AI-empowered software.

"Building AI-empowered software will demand a new breed of software professional, the AI engineer," said Walsh. "The AI engineer possesses a unique combination of skills in software engineering, data science and AI/machine learning (ML), skills that are sought after."

According to a Gartner survey conducted in the fourth quarter of 2023 among 300 U.S. and U.K. organizations, 56% of software engineering leaders rated AI/machine learning (ML) engineer as the most in-demand role for 2024, and they rated applying AI/ML to applications as the biggest skills gap.

To support AI engineers, organizations will need to invest in AI developer platforms. AI developer platforms will help organizations build AI capabilities more efficiently and integrate AI into enterprise solutions at scale.

"This investment will require organizations to upskill data engineering and platform engineering teams to adopt tools and processes that drive continuous integration and development for AI artifacts," said Walsh.

COOLING | HEATING | VENTILATION | CONTROLS

Ask ME about our IT Cooling solutions for Data Centres

Multi Density

Scan here

Data centres are now a crucial part of modern life and Mitsubishi Electric's **Multi Density** system makes IT Cooling easier than ever. Combining the efficiency of VRF with high-performance close-coupled air conditioning, Multi Density offers precise temperature and humidity control for high-density applications. It ensures optimal conditions for your IT equipment whilst helping increase energy efficiency and lowering emissions.

With features like advanced inverter technology, a small footprint, and pipe runs up to 165m, Multi Density is perfect for cooling high-density racks and blade servers.

Discover more about our **Data Centre Solutions** at: **les.mitsubishielectric.co.uk**

MITSUBISHI ELECTRIC

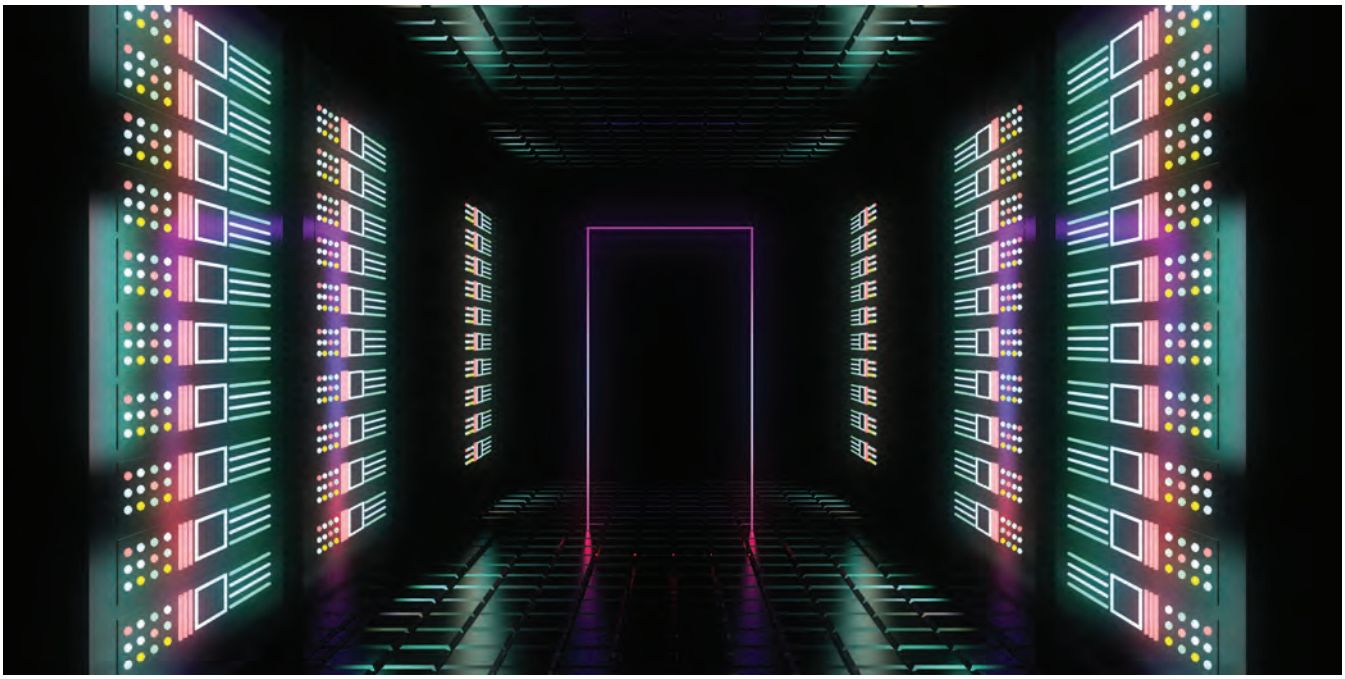

Where precision meets density:

unlocking the future of data centre testing

As data centres evolve to meet the demands of high-speed data transmission, the role of optical devices becomes increasingly critical. The efficient conversion of electrical and optical signals necessitates precision testing solutions, particularly for highly integrated optical devices. While there are profound challenges in testing these devices, innovative solutions exist to enhance testing efficiency and accuracy.

BY GOBINATH TAMIL VANAN, PRODUCT MANAGER, KEYSIGHT TECHNOLOGIES

AS DATA CENTRES ADVANCE to meet the demands of high-speed data transmission, the pivotal role of optical devices becomes increasingly evident. The efficient conversion of electrical and optical signals requires precision testing solutions, particularly for highly integrated optical devices. Optical devices play a fundamental role in modern data centres by efficiently converting electrical and optical signals. Increasing demands for higher speeds, smaller sizes, and improved data traffic drive the development of highly integrated optical devices. The advanced devices integrate multiple functions and components into a single unit, facilitating efficient system miniaturisation.

Testing highly integrated optical devices requires a significant number of precision bias sources. For example, as shown in Figure 1, testing integrated tunable wavelength lasers requires precision current sources for the laser diodes to ensure stable optical performance. It also requires precision bias sources for each heater to adjust the wavelength precisely towards the semiconductor optical amplifier (SOA). Similarly, coherent optical transceivers need multiple precision bias sources synchronised precisely to the phase control electrodes to convert electrical signals to optical signals accurately.

A detailed characterisation with very fine bias sweeping steps is necessary to test the optical power and wavelength of tunable lasers and coherent receivers. As a result, there are issues such as significant extension of test time and unintended wavelength shifts due to thermal effects. One effective solution is minimising each sweep step's duration, enabling faster sweeping to help mitigate these issues.

The need for optical components testing

Optical components are the lifeblood of high-speed data transmission systems. From receivers

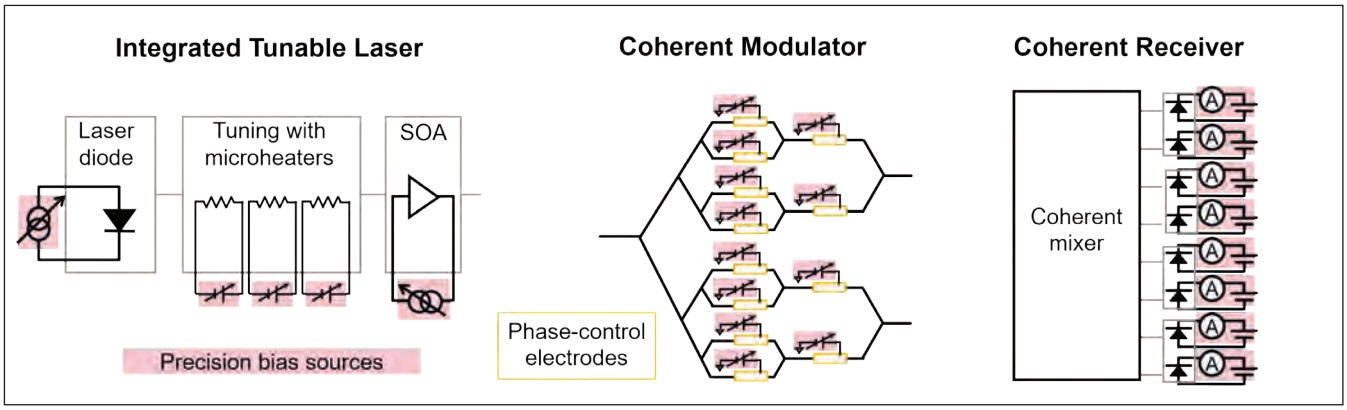

converting optical signals into electrical signals to modulators shaping and encoding data onto optical carriers, these components form the backbone of modern communication networks.

The relentless pursuit of higher data transmission rates and the demand for greater bandwidth in data centres place an immense burden on these optical components. The critical significance of testing these components lies in ensuring their reliability, performance, and compatibility within the dynamic environments of data centres.

Optical component testing becomes a key in guaranteeing that the components can withstand the rigors of continuous operation. Rigorous testing helps identify potential weaknesses, vulnerabilities, or performance limitations in these components, enabling engineers to refine designs and implement improvements. As data centre architectures evolve and the demand for energy efficiency increases, testing becomes instrumental in maximising optical devices' power consumption and thermal characteristics.

Accurate testing also helps validate theoretical models and simulations, ensuring that the behaviour of these components in real-world scenarios aligns with expectations. Engineers must be confident that the optical components will perform reliably under various conditions, including temperature variations, power fluctuations, and signal distortions.

In addition to ensuring the robustness of individual components, testing plays a crucial role in the overall system integration. It makes sense to identify and address compatibility issues, signal integrity challenges, and interoperability concerns during the testing phase to prevent potential pitfalls in the later stages of data centre deployment.

Compliance with industry standards and regulations is paramount; thorough testing is the key to meeting and exceeding these benchmarks. Whether it is the

➤ Figure 1. Integrated optical device testing requires numerous precision bias sources

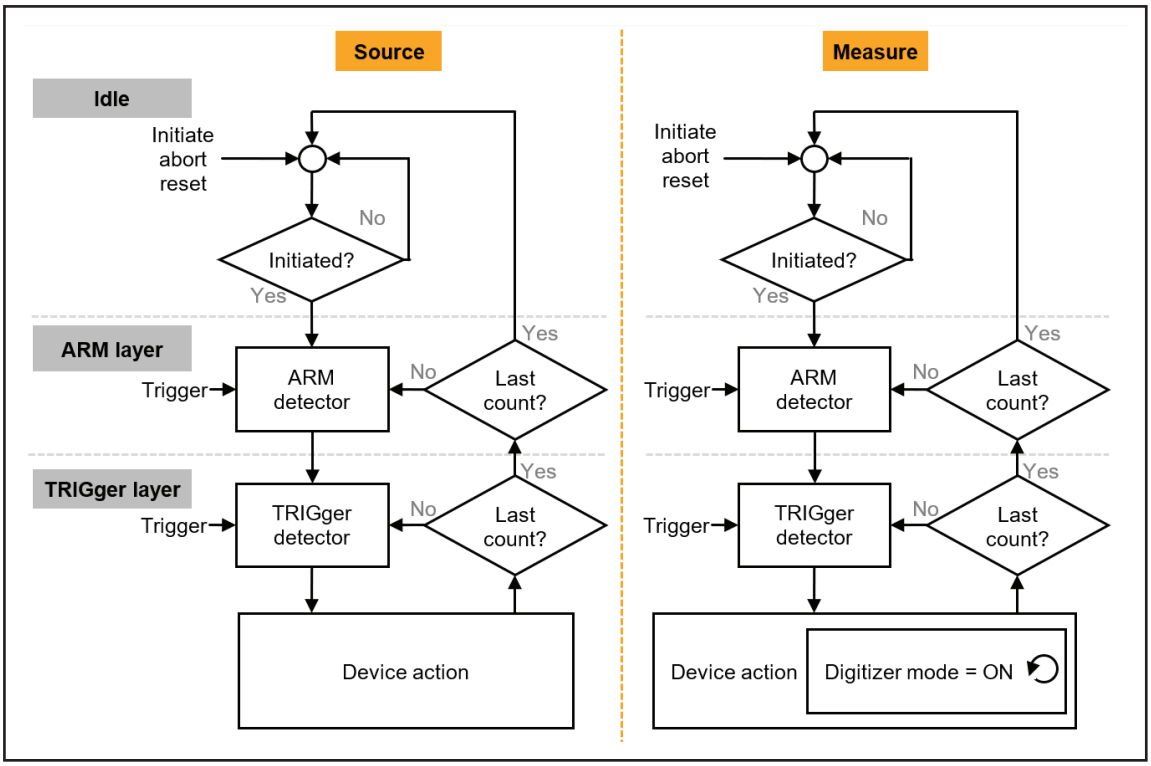

➤ Figure 2. Flow chart of an intelligent trigger system for flexible synchronisation

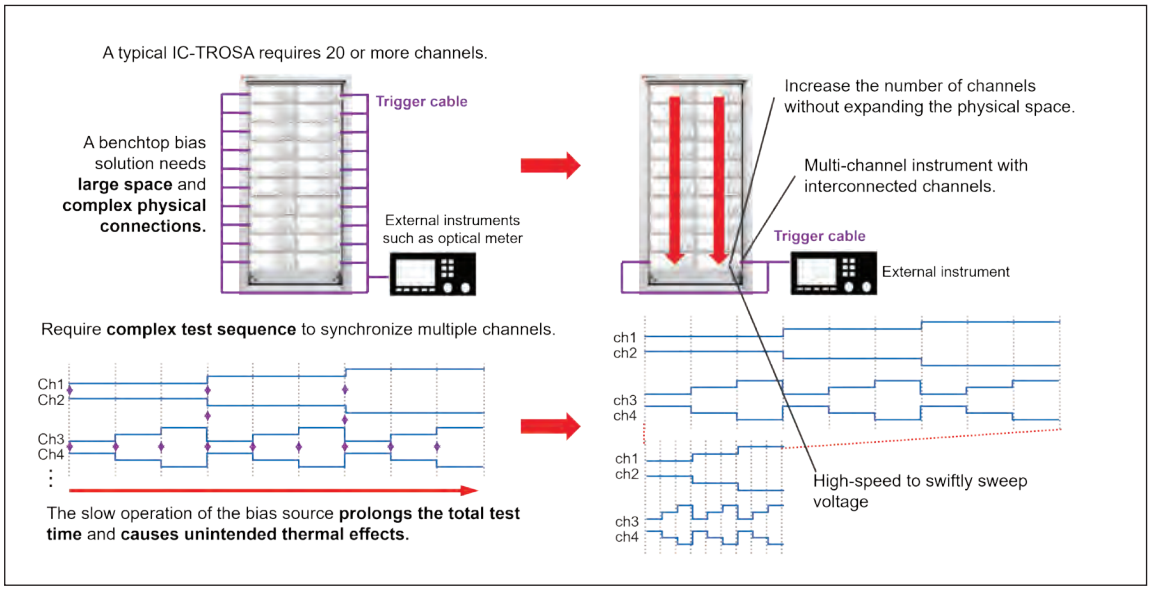

Figure 3. Examples of a 20-channel bias solution setup

standards related to optical power levels, signal-to-noise ratios, or bit error rates, testing ensures that optical components adhere to the stringent criteria set by the industry.

In essence, the importance of optical component testing extends far beyond mere quality control — it becomes a strategic step in advancing the capabilities of data centre networks. By subjecting these components to rigorous testing protocols, engineers pave the way for innovations that drive the data centre infrastructure's efficiency, reliability, and overall performance into the future.

The landscape of optical device testing challenges

Navigating the complex landscape of direct current (DC) bias testing for optical components in a data centre environment creates several challenges for engineers. Let's dissect these challenges to understand the hurdles engineers face:

Precision Requirements: The precision required to control bias voltage and current is a significant challenge. Optical devices, inherently sensitive to variations in biasing, demand a level of precision that pushes the limits of conventional testing equipment. Achieving and maintaining the necessary precision is challenging due to optical components' low tolerances and dynamic nature.

Dynamic Operating Conditions: Data centre environments are dynamic, with fluctuations in temperature, power, and signal conditions being the norm rather than the exception. Maintaining a stable DC bias under these dynamic operating conditions presents a major challenge. Optical components must operate reliably and consistently, even when subjected to rapid changes in bias levels, and the need to test these is paramount.

Non-Linear Behaviour of Modulators: Modulators, a critical component in optical communication systems, exhibit non-linear behaviour that complicates the testing process. Traditional testing equipment may need help to capture and replicate the intricate modulation characteristics accurately, which can lead to potential inaccuracies in assessing the performance of modulators under realistic operating conditions.

Sensitivity of Receivers: Optical receivers, responsible for converting optical signals into electrical signals, are susceptible to variations in bias levels. Achieving a stable and accurate bias for receivers is a delicate task, as even slight deviations can impact signal quality and, consequently, the

Data centre environments are dynamic, with fluctuations in temperature, power, and signal conditions being the norm rather than the exception. Maintaining a stable DC bias under these dynamic operating conditions presents a major challenge. Optical components must operate reliably and consistently, even when subjected to rapid changes in bias levels, and the need to test these is paramount

reliability of the entire communication system. Accurately capturing the significant current variations corresponding to light will also be extremely challenging.

Increasing Channel Density: Highly integrated optical devices have more test ports and components, requiring numerous high-precision power supplies and significant space. For instance, integrated tunable lasers need precision current sources for the laser diodes to ensure stable optical performance and precision bias sources for each heater to adjust wavelength precisely. Coherent optical modulators require multiple precision bias sources synchronised precisely to the phase control electrodes to convert electrical signals to optical signals accurately.

Real-World Simulation: Simulating real-world scenarios in the lab environment is a challenge. Engineers must ensure that the testing conditions accurately reflect the complexities of data centre operations. This includes simulating the varying conditions optical components may encounter in a live data centre, from load changes to ambient temperature fluctuations.

In summary, there are multiple challenges in DC bias testing for optical components in data centres: density, precision, dynamic conditions, non-linear behaviours, sensitivity, high-speed demands, and realistic simulations. Addressing these challenges requires innovative approaches and specialised equipment, making the role of source measure units (SMUs) crucial in overcoming these complex hurdles.

Testing optical devices with precision and density

To overcome the multifaceted challenges in DC bias testing for optical components in data centre environments, engineers turn to versatile SMUs. Let's delve into specific details highlighting why SMUs are crucial in addressing each challenge.

Precision and Stability

Achieving precision in bias voltage and current is where SMUs truly shine. SMUs come with ultra-high precision capabilities, enabling engineers to set and maintain bias levels with unprecedented accuracy. SMUs provide excellent stability, ensuring optical components receive consistent and reliable bias conditions. With low-noise DC signals, SMUs mitigate the risk of introducing unwanted interference that could compromise the accuracy of test results.

Intelligent trigger control

SMUs excel in dynamic bias control, a critical feature when dealing with optical components operating in dynamic data centre conditions. Some SMUs have additional capabilities, such as an intelligent trigger system for high-speed timing controls, as shown in Figure 2. The dynamic capabilities of SMUs enable engineers to simulate rapid changes in bias levels, replicating the real-world scenarios optical components face in high-speed data transmission environments. This not only ensures the accuracy of testing but also provides insights into how optical components perform under dynamic operating conditions.

Non-Linear Behaviour of Modulators

SMUs manage the non-linear behaviour of modulators. The programmability and precision of SMUs enable engineers to capture and reproduce modulators' modulation characteristics accurately. By providing a stable and controlled bias environment, SMUs allow for in-depth analysis of modulator performance, ensuring that testing results align with real-world expectations.

Sensitivity of Receivers

Addressing the sensitivity of receivers is a forte of SMUs because they offer the fine-tuned control necessary to provide stable bias conditions for receivers. With SMUs, engineers can tailor the bias parameters to match the sensitivity of optical receivers, ensuring precision over wide-ranging currents and repeatable testing. The precision of SMUs becomes particularly critical in scenarios where even slight deviations in bias levels can impact the performance of optical receivers.
High-Density Compact Form-Factor

With a high-channel density SMU form factor, as shown in Figure 3, you can save valuable rack space and minimise the test system footprint. Some flexible SMUs allow for any mixed module configuration for flexible scalability. An all-in-one SMU with integrated pulsar and digitiser functionality can reduce the required test instruments and system footprint. The SMU will address the challenges of optical device testing by providing a multi-channel, high-precision bias in a small footprint. The precision and easy integration capabilities streamline the evaluation process for optical device testing, saving significant space and improving test efficiency.

Real-World Simulation

SMUs facilitate realistic simulations in lab environments. Engineers can replicate the diverse conditions in live data centres thanks to the dynamic and programmable nature of SMUs. Whether it is simulating changes in load, fluctuations in ambient temperature, or other dynamic factors, SMUs provide the flexibility to ensure that optical components are tested under conditions that closely mimic real-world scenarios.

Conclusion

SMUs are indispensable tools in addressing the density, precision, intelligent trigger control, non-linear behaviours, sensitivity, high-speed demands, and real-world simulation challenges in DC bias testing for optical components. Their versatility and precision make SMUs essential for reliable, high-performance optical components in data centre environments.

New technologies, new frontiers

Revolutionising data centre maintenance for the Digital Age

BY ANNA MAZZOLENI, GLOBAL PRODUCT MANAGER, ELECTRIFICATION SERVICE, ABB

IN 2023, it was estimated that there were now 5.4 billion people online globally. With more of our lives increasingly being lived virtually, businesses, governments and consumers are now undeniably reliant on digital services. Forming the foundation of these are droves and droves of data centers which have become an integral part of our data-driven society, making the stakes for their operators higher than ever.

Cybersecurity risks notwithstanding, having a firm grasp on inherent operational and structural risks is critical. In the context of data centers, this means being able to provide a constant and reliable stream of data services. Yet, the threat of power failures that cause data center outages remains a significant challenge — and the resulting downtime that comes with it is not only highly disruptive, but costly.

With our dependence on data centers showing no signs of abating, we need to get ahead of these challenges. For one, enhanced approaches to asset servicing and maintenance through the use of innovative technologies, such as augmented reality (AR), are a way to mitigate these risks. Together, these offer a promising path forward for safer, smarter and more reliable data center operations.

Enabling the next frontier in servicing innovation
It's well-chronicled across a myriad of industries that the pandemic catalyzed novel ways of doing business — asset servicing and maintenance are no exception to this, and we've seen this specifically through the use of AR.

With an AR-enhanced maintenance app on their smartphones, technicians can see more than what meets the eye: technical information and servicing guidance through images, instructional videos and documentation are overlaid on the physical equipment in front of them.

This is done in such a way that enhances their perception and interaction with their surroundings, all while preserving the necessary level of environmental awareness to ensure the safest user experience while minimizing human error.

Such maintenance apps can also be designed to be device-agnostic, accessible across mobile, tablet and desktop, but also hands-free devices such as industrial smart glasses. This would enable engineers and technicians to operate hands-free, improving both safety and productivity as they conduct their assessments and repairs in the harshest environments.

Scaling servicing capabilities

Through AR, technicians also benefit from real-time, remote support from experts, assisting them with complex repairs and maintenance tasks. For data centers, especially, where every second of downtime translates to lost revenue, remote support is especially vital as it provides access to near instantaneous expert support and solutions to the technical issue at hand. This enhances first-time fix rates while also extending service reach to remote and under-served areas.

With data centers often located in remote locations, the benefits as it pertains to cost- and time-efficiencies are clear. But these equally extend to sustainability considerations: Remote servicing saves up to 171g of CO2 emissions per passenger-kilometer, totaling to 332 tons of CO2 emissions per year by reducing up to a third of customer site visits by field service engineers.

At the same time, the use of technology to power remote support capabilities equally translates to more opportunities for immersive remote training which is proven to deliver better knowledge acquisition and retention. This, compounded by self-learning on virtual systems fosters a culture of first-rate support and servicing. AR can also power hands-on training experiences, enabling technical servicing teams to scale their junior talent. Junior technicians benefit from practical training opportunities through a virtual system before applying their skills in the field.

Taking the smart approach to maintenance

Businesses tend to maximize their capital expenditure by running their equipment, unknowingly, to the point of failure, which leads to longer term losses. In fact, this can cost up to ten times more than investing in regular maintenance and often contribute to more severe outages — and the older the equipment is, the greater the severity. According to the Uptime Institute's 2023 Global Data Center Survey, such outages can cost from US$250,000 to more than US$1 million.

Regular monitoring of facilities, especially mission-critical equipment, is essential for guaranteeing reliable service provision — and smart maintenance can better support that.

A preventative approach involves replacing older, non-digital circuit breakers with intelligent, sensor-enabled breakers linked to cloud-computing

platforms. Real-time data and analysis on asset condition and performance can prevent potential issues before they arise.

According to the Deloitte Analytics Institute, predictive maintenance increases productivity by 25%, reduces breakdowns by 70%, and lowers maintenance costs by 25%. Despite these benefits, less than 50% of global manufacturers use predictive maintenance technologies. The gap itself presents an opportunity for businesses to increase productivity and build a competitive advantage.

Digitally driven, digitally resilient

For businesses today, the incorporation of more digitally-integrated approaches that can ensure greater reliability is no longer a luxury but a necessity — and it's simply no different for data centers. With the cost of outages increasing and data center downtime costing nearly US$8,000 per minute, every second of downtime translates to revenue lost and reputational damage that's hard to come back from.

As the backbone of our increasingly digital world, data centers are pivotal in ensuring the smooth operation of critical services, be it in healthcare or financial services. To keep them running, we need to rethink our approach to servicing and maintenance, bringing outdated methods into the future. Powerful innovations such as augmented reality unlock new opportunities for predictive, data-driven maintenance all while also significantly reducing environmental impact and operational costs.

Ultimately, taking a more digitally-enabled approach to asset maintenance is essential to navigating the challenges associated with increased digitalization. Thankfully, innovative technologies are set to play a pivotal part in ensuring data centers to remain resilient and operationally-ready to meet the growing demands of data services.

Data centres: The good they do

The positive impact of data centres on people, society, business and government.

BY ED ANSETT, CHAIRMAN AND CO-FOUNDER, I3 SOLUTIONS GROUP, PART OF RAMBOLL

DATA CENTRES continue to evolve, driving innovation and transforming how we live, travel, and enjoy our leisure time.

Data centres are the unseen 21st century utility with which people interact hundreds of times daily – mostly without any awareness that their financial transactions, communication, shopping, leisure and travel are passing through physical data centre infrastructure.

Smart phone apps are becoming the tools around which people build their routines for work, play, family interaction and an increasing number of 'life management' activities.

Social media is one of the most obvious cases in point. Each interaction, be it a 'like' a post or a message creates data uploads and downloads.

The sheer volume of interactions tells it own story. Backlinko statistics show that Facebook has 3.03 billion monthly active users, with YouTube receiving 2.49 billion, WhatsApp and Instagram 2 billion each, and Wechat 1.33 billion.

In the gaming world, it has been speculated (by Elait) that a single popular online gaming platform generates around 50 terabytes of gaming data due to its extensive user base.

According to Statista, the anticipated user penetration rate for video streaming services is forecasted to be 18.3% in 2024, reaching 20.7% by 2027.

Travel has been transformed by the services enabled by data centres such as websites and apps designed streamline booking (whether flights or accommodation) and make trip-planning easier and more efficient.

For example, booking.com amassed approximately 556 million visits in 2024, with tripadvisor.com and airbnb.com receiving roughly 176 million and 101 million users respectively, as reported by Statista. Beyond the booking of trips and holidays, travellers have greatly gained from GPS and its related apps, which require large amounts of data centre storage and processing power. A National Institute of Standards and Technology-sponsored study estimates that GPS has generated around $1.7 trillion in economic benefits for the USA. What's more, digital translation tools and apps bridge language barriers, facilitating communication in diverse travel settings.

Businesses and consumers have both benefited from the increasing popularity of online reviews and recommendations.

The popularity of fitness apps soared with a user base of 368 million in 2023 with the apps were downloaded over 850 million times. Fitness apps and wearable devices track physical activity, provide workout routines, and promote overall wellness through personalized guidance and monitoring. According to Business of Apps statistics the fitness app industry saw a significant surge, generating a revenue of $3.58 billion in 2023 thus marking a 9.1% increase from the previous year.

What many of these services have in common is what is commonly known as 'The Cloud.' According to a recent study by the Institute for Business Value (IBV), 97% of enterprises have either piloted or integrated cloud into their operations.

How data centres benefit society

Data centres are fundamental to the digital infrastructure that powers modern life, from daily communication and entertainment to critical services like healthcare and finance.

Tools such as video conferencing, cloud storage and collaboration software make remote-work setups possible, allowing for flexible hours and a healthier work-life balance. This helps the company as well as the individual: a Connect Solutions survey found that 77% of individuals who work remotely at least a few times each month show improved productivity.

Data centres also help make people more mobile. According to a report from Ericsson, the average mobile data usage per smartphone is forecasted to rise from 21 GB in 2023 to 56 GB by 2029. Another report by Ericsson suggests that Fixed Wireless Access (FWA) is expected to reach 130 EB per month at the start of 2024, with a projected 403 EB growth per month in 2029. It is data centres that enable the instant communication, via email and messaging apps which we have come to rely on in our daily lives.

Our shopping, too, is increasingly reliant of data centres: the e-commerce market is expected to be worth $6.3 trillion by the start of 2024, with China leading in online retail sales globally. Online sales are predicted to comprise over 21.2% of total retail sales in 2024, with studies in 2023 indicating that 218.8 million US residents regularly made online purchases, according to a report by Artois.

It is obvious that the smartphone has changed society. What is less obvious is that it can only operate as the end point device of a vast array of networked data centres.

What data centres do for business

Data centres also facilitate data analysis, meaning data-driven decision-making and better business. But the digital tools powered by data centres are not just a matter of corporate streamlining.

Data centres ensure the security and integrity of financial data. A study by the Boston Consulting Group found that 43% of customers are willing to switch banks due to a disappointing digital experience: this highlights the importance for fintech companies of using modified technology.

Across the global financial markets high-frequency trading, algorithmic trading, risk analysis, and other computationally intensive financial operations rely on the support of data centres, which also help financial institutions comply with regulatory requirements such as GDPR (General Data Protection Regulation).

Work:

- **Digital Banking:** 30% of banks have undergone a digital transformation report successfully implementing their digital strategy according to a report from McKinsey.
- **Data Analysis:** Digital tools enable organizations to gather and analyse data, leading to data-driven decision-making and improved business outcomes. Statista predicts by 2025 global data creation is projected to grow to 180 zettabytes.
- **Global Reach:** Digital platforms enable businesses to reach global audiences, expanding market opportunities and facilitating international collaborations

Healthcare is becoming increasingly digital and so increasingly relying on data centres. They enable the secure storage and analysis of health records and imaging data such as MRI scans and X-rays. Data centres are integral to a hospital's infrastructure as well as day-to-day management: the average hospital generates 50 petabytes of data annually, with healthcare data increasing by 50% per year (according to Becker's hospital review).

As for individual healthcare, it is the telemedicine services, which depend on data centres, that enable remote consultations with healthcare professionals, improving access to medical advice and reducing healthcare barriers. Yahoo Finance has projected the value of the global healthcare predictive

analytics market is set to hit $34.1 billion by 2030, growing at a compound annual growth rate (CAGR) of 20.4% from 2024 to 2030. The UK's NHS App has surged to nearly 28 million users, marking a substantial increase from approximately 2.2 million users in September 2020.

Data centres are also key to the transportation industry, whether supporting the real-time tracking and monitoring of vehicles in delivery services or processing and analysing data from traffic sensors, cameras and GPS devices to reduce congestion and improve traffic safety. This will play an important role in autonomous vehicle innovation, a market which will amount to $27 billion by 2025, according to Omdia.

Autonomous vehicles will rely on data centres for the analysis of sensory data in predictive maintenance: NVIDIA DRIVE, for example, is a system for the creation of secure autonomous vehicles which combines data centre hardware and software.

Digital media content is delivered with low latency and high reliability to its users worldwide thanks to the content delivery networks (CDNs) supported by data centres. The vast amounts of user-generated data from social media and streaming platforms which is processed by data centres provides content creators and marketing strategists with useful information about user behaviour, preferences and trends. International Data Corporation (IDC) predicts that this user-generated data will continue to grow at an annual rate of roughly 21% and exceed 221 Zettabytes by 2026.

IDC network functions virtualisation infrastructure (NFVI) will grow from $12.9 billion in 2022 to $27.3 billion in 2027—and, in hosting telecommunications equipment, including servers, switches, and routers, data centres are key to supporting those voice, data and video communication services. Data centres also support Edge Computing solutions for low-latency application as well as the rollout of 5G networks by providing centralised data processing and storage. According to Business wire, North America will host over 335,000 5G small cells as Communications Service Providers (CSPs) by 2025. Disaster recovery, on which telecom companies spend 2-4% of their telecom budget (as estimated by Numeric Futures), is another domain of data centre support for the telecommunication industry: data centres ensure the continuation of telecommunication services through backup systems and recovery plans which minimise downtime and maintain service reliability.

The most high profile impact of data centres on business is the AI revolution. Much has been written about AI and data centres and energy requirement forecasts. The AI software market is going from strength to strength. In 2022 the AI software sector was calculated to be worth $328 billion. According to a report from blend by 2029, the market is expected to be worth approximately $1.3 trillion – almost as much as the entire IT industry is today.

What data centres do for government

Governments across the world are using data centres and cloud technologies to modernise services and meet the evolving needs of citizens and businesses.

The US, for example, has embraced cloud computing through the initiative Cloud First Strategy, which encourages agencies to prioritise cloud solutions for IT deployments. A similar initiative, the European Cloud Initiative was launched by the European Commission to promote cloud adoption, data sharing and digital innovation across EU member states.

Through the UK's Cloud First Policy, several UK government departments have moved their services to cloud platforms (such as Amazon Web Services, Microsoft Azure, and Google Cloud Platform) to improve scalability and efficiency. India's Digital India campaign aims to do the same by transforming government services through cloud computing, data analytics and e-governance platforms.

Japan's Smart Cities Mission, focus on leveraging data analytics and the Internet of Things to improve citizen services and keep urban development sustainable. Digitalisation of government services is increasing and accelerating.

Conclusion

Data centres have entered the public consciousness. Unfortunately for the sector this has mostly been through media coverage that, with some justification, primarily focuses on the cost in energy and natural resources required to build and operate this new breed of 21st century utility infrastructure.

And just as increases in demand do not absolve the energy industry of responsibility for providing clean, sustainable power (with which data centres are intrinsically linked) so the DC industry must face the demand curve for digital services as a challenge by taking responsibility for communicating its benefits while accepting and continually tackling and reducing its environmental cost.

wesco | anixter

CHATSWORTH PRODUCTS

Protect

Global Expertise in Protective Data Centre Solutions

Delivering value at every stage of data centre development and operations, Wesco Anixter brings reliable expertise alongside a global ecosystem of strategic partners. Together with CPI, we provide advanced rack and cabinet solutions to protect your cabling infrastructure, ensuring robust security and reliable connectivity to keep your network running efficiently.

Scan to learn more about CPI's rack and cabinet solutions.

We build, connect, power and protect the world.

Wesco.com/datacenters-emea

How AI can slash data centre costs

AI can help to improve the basic functionality of data centres, enabling businesses to ensure their infrastructure is fit to serve the next great tech boom - and to reduce their carbon footprint.

BY SUJATHA IYER, MANAGER OF AI IN SECURITY, MANAGEENGINE

THE GENERATIVE AI boom has seen demand for data centre power soar. Capacity demands had already been increasing rapidly due to the near-universal adoption of cloud-based apps and storage across businesses and for private use cases. Now, advanced AI models have burst into the scene in a similar way, bringing with them a need for unprecedented computing power. At its heart, AI is based on the rapid analysis of vast amounts of data—which means data centre capacity is having to grow exponentially to keep pace.

And because of all this, power needs are jumping, too. Driven in part by AI adoption, data centres across the globe are consuming ever-increasing amounts of electricity. Two years ago, data centres consumed 460 terawatt hours, and the International Energy Agency predicted this figure was set to double within four years. That means that by 2026, data centres could be using 1,000 terawatt hours annually. What does that mean? Well, it's roughly equivalent to the electricity consumption of Japan, which has a population of 125 million people. Our technology needs are effectively threatening to add a whole new developed nation's worth of energy usage to the global total.

There are clear issues with this—most pressingly, the ecological and environmental impact this energy use could have. The electricity that powers data centre racks has to come from somewhere, and the amount that's generated via clean or renewable sources varies wildly across the world. As AI usage grows, the IT infrastructure industry needs to play its part in delivering solutions that work for the planet as well as end users. At the same time, cost is the other significant consideration; if generative AI is to continue flourishing, it's important that the associated costs, including energy use, are carefully managed.

AI: Cause or cure?

But the cause of the problem could be the solution. As well as drawing on resources, AI can transform data centre usage. Using AI in data centres could lessen energy consumption, improve efficiency, and provide financial savings. For example, AI can help keep data centres cool by analysing historical data and predicting when systems will face increased workloads and, inevitably, the high temperatures that come with them.

AI is highly effective at optimising processes, so why not use it to optimise data centres? With the help of AI, data centre operations can be optimised by putting operational data to work

Additionally, AI requires a lot of electricity to operate. Naturally, AI processing in data centres can result in higher heat output, which seems to be a counterproductive approach when implementing AI for sustainability purposes, especially if it is to be used for servers and model training. This explains why the majority of AI developments in data centres are concentrated on simple functions, such as enhancing rack-level cooling and power efficiency.

Indeed, many of the existing problems that companies are facing with traditional data centres, such as thermal management and data breach and security issues, can be eliminated through AI use—allowing the industry both to respond to the emerging challenge of growing energy use and to streamline operations, ultimately enabling new growth.

Let's look in more detail at two key ways AI can contribute to a more efficient data centre estate: workload management and predictive analytics.

Workload management

AI is highly effective at optimising processes, so why not use it to optimise data centres? With the help of AI, data centre operations can be optimised by putting operational data to work. With the right intelligent systems in place, data centre resources can be allocated in a way that reduces downtime and minimises energy wastage. AI can analyse energy-intensive and routine tasks, dynamically utilising energy based on the processes.

At the same time, the development of GPU-based AI data processing and LLMs means fundamental performance can be significantly enhanced, achieving up to 50 times the processing speed of CPU-based systems. This advancement means that a few hundred GPU-based systems could replace tens of thousands of CPU-based servers, leading to reduced server space, lower energy costs, decreased core-based software licensing fees, and simplified management.

There are many other processes that contribute to the effective functioning of a data centre, like maintaining the temperature of the entire system. Data centres generate a lot of heat due to the workloads they handle. If the machines in the data centre aren't kept cool, they're liable to overheat, break down, and bring operations to a halt. So you're not just using energy to run the racks—you're also using (a lot more) energy to cool them.

Here, AI can lend a hand by analysing the system and more effectively utilising the power available to keep the data centre temperature controlled and cool.

Predictive analysis

Another domain where AI greatly helps is in the analysis of data. As data centres become more complex, it is increasingly important to analyse historical data to predict factors like workload and temperature. These systems improve over time, and as AI advances, they can analyse more data and predict these factors with greater accuracy.

Data centre companies are already using machine learning to create intelligent thermal management solutions, aiming to reduce energy consumption in data centre infrastructure. These kinds of predictive systems can achieve a significant reduction in power usage, leading to significant annual power cost savings of many hundreds of thousands of kilowatt hours—which translates to millions of kilograms of carbon dioxide.

In essence, AI is at its best when it's automating tasks that would take a human hours or days of laborious, repetitive work. By using its formidable capabilities to improve the basic functionality of data centres, businesses can ensure their infrastructure is fit to serve the next great tech boom—and reduce their carbon footprint.

It’s time to rethink data centre power

The AI revolution is fully under way. But without a top-to-bottom reassessment of power solutions, its progress could slow to a crawl.

BY TOD HIGINBOTHAM, COO, ZINCFIVE

WITH BUSINESSES clamoring to harness the potential of AI, data center operators have a mandate to outfit their facilities with the latest CPUs, GPUs and other components that power the most demanding workloads. That level of power, however, has sparked a surge in electricity needs. And meeting those needs is harder than it sounds.

Already, data center projects are being hamstrung by power problems. Data center operators need more power -- and they need it from clean power sources. They also need a transmission system that can handle a heavier load. The burden, however, doesn’t belong solely to utility providers. Data center operators also need to innovate. By rethinking data center design, operators can minimize their power needs. Meanwhile, they need to strategically monitor and manage their power usage to ensure optimal operations.

The AI revolution is fully under way. But without a top-to-bottom reassessment of power solutions, its progress could slow to a crawl.

A surge in demand

Insufficient power resources are already slowing projects down by years, according to a recent report from Cushman & Wakefield, a global commercial real estate services firm.

“Over the past year, power has become the number one consideration for data center operators as they conduct site selection to rapidly grow their portfolios,” the firm’s 2024 Global Data Center Market Comparison says. “Many utility providers are suggesting wait times of 2-3 years or more for sizable power to be delivered to their developments.”

Data center operators are expanding their footprint in anticipation of the growth of AI. Currently, hyperscalers need around 10 kW to 14 kW per rack, the commercial property consultancy Newmark noted in a recent report. However, AI workloads will push that requirement up to 40 kW to 60 kW per rack. All told, Newmark expects AI to drive US data center demand to 35 GW by 2030, up from 17 GW in 2022.

As they anticipate greater power demands driven by AI, data center architects should consider what types of AI loads they are building for -- in other words, whether a data center operator will manage primarily inference or training. With AI training, there

DCS AWARDS 2025

CELEBRATING 15 YEARS OF SUCCESS

The DCS Awards: 38 Categories across 5 Themes

22 MAY
Leonardo Royal Hotel
London St Pauls

KEY DATES:

14 FEBRUARY: NOMINATIONS CLOSE

21 MARCH: SHORTLIST ANNOUNCED

24 MARCH: VOTING OPENS

23 APRIL: VOTING CLOSES

22 MAY: AWARDS CEREMONY

NOMINATE NOW!

SPONSORSHIP PACKAGES

The DCS Awards offer extensive branding and sponsorship opportunities through online advertising in our Datacentre Solutions & Digitalisation World publications, and of course at the awards ceremony itself.

For sponsorship opportunities and/or to book your awards table please contact: awards@dcsawards.com or call +44 (0)2476 718970

NOMINATE: https://dcsawards.com/nominate

Supported by

The data centre trade association

may be extreme load steps that require larger utility feeds or the use of innovative battery storage. There is interest in using batteries to offload these peaks and spare the uninterruptible power supply (UPS) from these load steps. These batteries could be located relatively close to AI servers and allow for a significant reduction in utility MWs to the facility.

Securing access to different power sources, transmission lines

When data center architects plot out new builds, they aren't just thinking about power availability but also the source of that power. Across the globe, more stringent regulatory environments as well as pressure from corporate stakeholders are driving data center operators to step up their reliance on renewable energy.

The demand for renewable power is reflected in forecasts for U.S. energy production: Solar installations will account for "almost all growth" in US power generation in 2024-2025, according to the U.S. Energy Information Administration (EIA). Utilities are opting for solar installations, Reuters notes, thanks to tax credits available from the 2022 Inflation Reduction Act.

While the shift to clean energy is happening, data center developers still need to think strategically about siting new facilities. Even with the increase in solar installations, the EIA said that by 2025, solar will still only account for 7% of power production in the US.

Along with power generation, data center developers need to consider transmission. On both of these matters, developers can collaborate directly with utility providers to ensure they can access the power they need.

The US electric grid faces capacity shortfalls, the North American Electric Reliability Corporation (NERC) warned in December, due to increases in demand and fossil fuel generators coming offline.

Delivering sustainable backup power

Transmission lines should also be assessed for reliability. Aging transmission lines, as well as transmission lines impacted by extreme weather events, can potentially cause costly outages. Power issues are consistently the most common cause of serious and severe data center outages, the Uptime Institute found in its annual outage analysis. More than half of the respondents surveyed by the Institute said their most recent significant, serious or severe outage cost more than $100,000.

The costly nature of outages underscores the importance of backup power systems. As with other power sources in the data center, developers face an imperative to make backup systems sustainable. Nickel-zinc (NiZn) batteries, an innovation in battery technology led by ZincFive, can power UPS systems more sustainably than traditional lead-acid batteries or lithium-ion batteries. In comparison to other chemistries, NiZn produces lower GHG emissions and offers a smaller water footprint and energy footprint. Specifically, NiZn batteries' lifetime greenhouse gas emissions are 4x lower than lead-acid and 6x lower than lithium-ion emissions. Nickel-zinc batteries use common, widely available, conflict-free materials. They're also highly recyclable.

Meanwhile, major corporations are exploring a range of alternative energy sources for backup generators. Microsoft, for instance, has been testing the viability of using large-format hydrogen fuel cells to supply data center backup power. Microsoft is also installing a "resiliency microgrid," which relies on renewable natural gas, for backup power at its San Jose, Calif. data center.

Dealing with heat

Data center developers and operators should also be looking for ways to more efficiently manage the heat generated by their infrastructure. Cooling systems are all the more critical as businesses adopt AI and HPC systems that emit greater levels of heat. Typically, cooling systems consume around 40% of a data center's power.

In addition to offering sustainable backup power, NiZn batteries allow data center operators to reduce the footprint of their cooling systems and other safety infrastructure. Nickel-zinc batteries exhibit no thermal runaway at the cell level and are thus non-flammable, unlike other UPS battery alternatives.

Conclusion

While power plays second fiddle to innovations in computational hardware and increasingly sophisticated AI workloads, it's a fundamental part of the data center. And as with other elements in the data center, it's rapidly evolving. To build a data center that can meet the demands of tomorrow's workloads, data center developers and operators should reconsider their power infrastructure, top to bottom.

Available for sales and rental

CRESTCHIC

THE **TRUSTED** NAME FOR **LOAD TESTING**

200kW Data Centre Loadbank: Built for Integrated Systems Testing

Simulate full load operation, power outages, and cooling system failures

Verify PDU operational readiness during commissioning

Closely mimic server heat generation to test cooling systems

Control up to 20 units from one controller

- Small footprint, stackable for easy storage
- AC load bank resistive-only (200kW)
- Voltage range: 400V or 415V
- 50/60 Hz Operating frequency range
- 3 Phase
- Forklift pockets for lifting
- Sheet steel construction with IP55 rating
- Low Delta T
- Low noise
- Internal or external auxiliary supply operation

www.crestchicloadbanks.com

sales.uk@crestchic.com | hires.uk@crestchic.com | service.uk@crestchic.com

Sales & Rental Offices in : UK, North America, Singapore, Germany, Belgium, France, China and UAE

Transforming data centres to meet AI's evolving demands

While 2023 marked a pivotal moment in recognising the vast potential of artificial intelligence, 2024 kicked of what is becoming a truly transformative period. AI's broad applications, ranging from machine learning and deep learning to natural language processing have seamlessly integrated into our everyday lives, revolutionising how we live, work and connect.

BY KAMLESH PATEL, VP DATA CENTER MARKET DEVELOPMENT AT COMMSCOPE

As AI's popularity soars, data centre managers and their teams are grappling with the challenge of managing, not only the surge of petabytes of data flooding their networks, but also the need for ultra-low latency. Additionally, they are attempting to tackle the increased power demands and higher fibre counts needed to support the advancements that come from supporting AI.

Similarly, the rise of artificial intelligence has caused a fundamental shift in data centre design, significantly impacting network infrastructure in areas such as cabling, connectivity, architecture, resilience and adaptability.

Here are the key challenges and opportunities that I believe come with cabling AI data centres, some best practices and tips for success.

The unstoppable surge in power demand

Regions that house data centres are experiencing a surge in power demand. In the Republic of Ireland for example, data centres now consume over 20% of the country's electricity, a significant increase from just 5% in 2015. Consequently, for the first time ever, there is no longer a guarantee that the power needed to support data centre operations can be reliably supplied.

Recently, the 'net zero' goals of major tech companies have been challenged by this increasing power demand, a direct consequence of AI and energy-hungry data centres. Google reported a 48% increase in its greenhouse gas emissions over the past five years, largely due to the growth of its data centres, while Microsoft's Scope 3

emissions have risen by over 30% since 2020. To strike a balance between enhancing sustainability and expanding capacity and performance, data centres will require support from their infrastructure technology partners.

Ultra low latency meets ultra high connectivity solutions

Because the models used to train and run AI consume significant processing capacity and are typically too much for a single machine to handle, processing these large AI models requires numerous interconnected GPUs distributed across multiple servers and racks. This presents a unique challenge for the cabling infrastructure that links everything together to keep data flowing.

For instance, GPU servers demand significantly higher connectivity between servers, but due to power and heat limitations, fewer servers can be housed per rack. As a result, AI data centres require more inter-rack cabling compared to traditional data centres. Each GPU server is linked to a switch within the same row or room, with these connections needing 400G and 800G speeds over distances that traditional copper cables like DACs, AECs or ACCs can't handle. Moreover, every server must also be connected to the switch fabric, storage, and out-of-band management.

In an ideal setup, GPU servers in an AI cluster would be close together, because AI and machine learning algorithms - like high-performance computing (HPC) - are highly sensitive to latency. It's estimated that 30% of the time spent running a large training model is due to network latency, while 70% is spent on compute time. To reduce latency, AI clusters strive to keep GPU servers in close proximity, with most links limited to 100 metres. However, not all data centres can place GPU server racks in the same row. These racks require over 40 kW to power a GPU server, far more than typical server racks, forcing traditional data centres to space them out accordingly.

Although extra space isn't feasible in the densely packed server rack layouts of modern data centres, managing the narrow, congested pathways and the added cabling complexities brought by AI is made possible through innovations like rollable ribbon fibre.

The innovative design allows for the installation of up to six 3,456 fibre cables within a single four-inch duct, providing more than double the density compared to traditionally packed fibres.
In the rollable ribbon fibre cable, the fibres are attached intermittently to form a loose web. This design makes the ribbon more flexible, allowing the fibres to flex with a degree of independence from one another. The fibres can now be "rolled" into a cylinder, making much better use of space when compared with flat ribbons.

While the cables are lighter and simplify handling and installation, their intermittent bonding enables installers to position the fibres naturally into a smaller cross-section making it perfect for splicing.

In an ideal setup, GPU servers in an AI cluster would be close together, because AI and machine learning algorithms - like high-performance computing (HPC) - are highly sensitive to latency

Data centre architecture of the future

Looking to the future, the value proposition for data centres will hinge on their extensive processing and storage capabilities and operators need to thoughtfully select the optical transceivers and fibre cables for their AI clusters.

In an AI cluster, the optics cost is primarily driven by the transceiver due to its short links. Transceivers that utilise parallel fibres are particularly beneficial because they eliminate the need for optical multiplexers and demultiplexers, which are typically required for wavelength division multiplexing (WDM). This results in reduced costs and lower power consumption for transceivers with parallel fibre.

Links up to 100 metres are supported by both singlemode and multimode fibre applications and advances such as silicon photonics have lowered the cost of singlemode transceivers.

In many AI clusters, active optical cables (AOCs) are used to interconnect GPUs spread over many servers and racks. These cables are usually designed for short distances and are commonly used with multimode fibre and VCSELs. The transmitters and receivers in an AOC may be the same as in analogous transceivers but are the castoffs. These components don't need to meet stringent interoperability requirements since they are only required to work with the specific unit attached to the other end of the cable. Additionally, since the optical connectors are not accessible to the installer, there is no need for specialised skills to clean and inspect fibre connectors.

Strategic planning for AI cluster cabling

In summary, data centres must evolve and adapt to meet the growing demands of artificial intelligence in business applications and customer service delivery. Infrastructure designers and planners must focus on improving efficiency, scalability, and sustainability. Key to these advancements is the upgrade of cabling systems, which will help reduce costs, energy usage, and installation times. By embracing these innovations, data centre facilities will be well-equipped to manage both current and future AI-driven workloads.

Waste heat and PUE improvement

Project study considering EED targets and impact to PUE improvements.

BY DAWID KROPIWNICKI, PRINCIPAL MECHANICAL ENGINEER AND NICK REMINGTON, TECHNICAL DIRECTOR, BOTH AT BLACK & WHITE ENGINEERING

THE QUEST to improve power usage effectiveness (PUE) has been one of the long-standing driving forces in data centre (DC) design, procurement and operation. Black & White (B&W) Engineering has concluded a recent study exploring how the utilisation of waste heat can further improve PUE. PUE is a standard efficiency metric for energy use in data centres. It is the ratio of total facility energy to IT equipment energy used in a DC.

The application of PUE for DC facility is multifaceted, including consideration of maximising IT yield, greater utilisation of available utility, operational expenditure forecasting and indirectly as a sustainability metric. Furthermore, legislative PUE targets are now in place in some locations.

The European Energy Efficiency Directive (EED), effective from May 2024 in Germany, now stipulates PUE targets as well as waste heat percentage contributions to be provided to a local district heat network operator (DHNO).

An R&D case study has considered the opportunity to utilise the waste heat to aid in meeting these targets. The considered project is a notional 54 MWIT DC in Germany, deploying high efficiency free-cooling air-cooled chillers. Waste heat is in the form of low grade heat and does not include for heat pump technology to provide temperature elevation.

EN 50600's definition of PUE does not consider the use of waste heat in determining a DC's PUE. However, in response to the EED's stipulated PUE compliance and percentage contribution of heat export requirement, this study has continued to assess this opportunity. Any application of heat export to improve PUE and therefore deviation from EN 50600 shall be discussed with the local authority having jurisdiction.

Under the EED, waste heat is required to be provided to a local DHNO, unless no such operator exists or the heat supplied is not required. A local DHNO's heat demand profile, (typically heating and

hot water) has a higher demand during the peak winter months, reducing in the off-peak months to a base load during the summer. Typically, summer months have low heating demand due to external conditions negating heating loads.

Conversely, the PUE profile through the year is seasonal and a function of the free-cooling operation of the air-cooled chiller, the largest contributor to PUE after the IT load.

The above shows the optimum period of the year where chiller compressor power consumption can be omitted through heat rejection to a DHNO, is simultaneously when DHNOs do not require the available waste heat. Until the above profiles have greater convergence, PUE improvements are limited. It should be noted, alternative annual applications such as swimming pools, greenhouses, aquaculture and industrial processes would increase this heat profile baseload.

Using the same notional project case study, B&W have assessed hypothetical percentage increases in the DHNO base load, artificially increasing the opportunity for heat rejection to the DHNO. From a practical, physical and commercial viability perspective, high percentage heat export values are more difficult to achieve in practice and are only considered here to test this PUE improvement hypothesis.

Upon analysing the results, increasing the DHNO baseload from 0% to 100% IT load yields up to 3.3% reduction in annualised PUE, while the impact to peak PUE is more substantial, at 21.2%. When the parasitic PUE component (or PUE overhead) is considered in isolation, increasing the DHNO baseload from 0% to 100% IT load, offers up to 17.9% in the annualised parasitic PUE and an impressive 63.5% in the peak parasitic PUE.

On average, for every 10% baseload waste heat increase, the parasitic annualised PUE improved by ~2%. These results account for the increase in annual pump power consumption to serve the DHNO.

Conclusion

The opportunity to improve PUE through exporting waste heat is possible, subject to local acceptance and deviation from EN 50600.

The PUE improvement is heavily reliant on the DHNO usage profiles. To fully realise this opportunity, DHNOs are required to accept more heat for a larger period of the year, minimising chiller compressor power consumption. Alternative applications for the waste heat could also be considered.

Although PUE improvement is available, the viability of higher percentage values of heat export is required from a physical and commercial perspective.

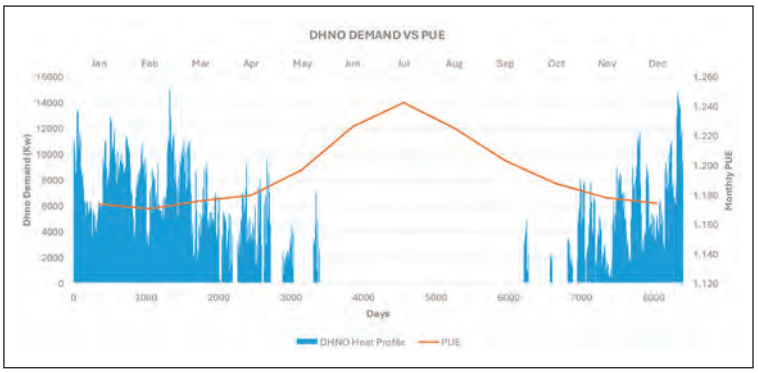

Figure 1: Typical Annual DHNO profile vs Monthly PUE

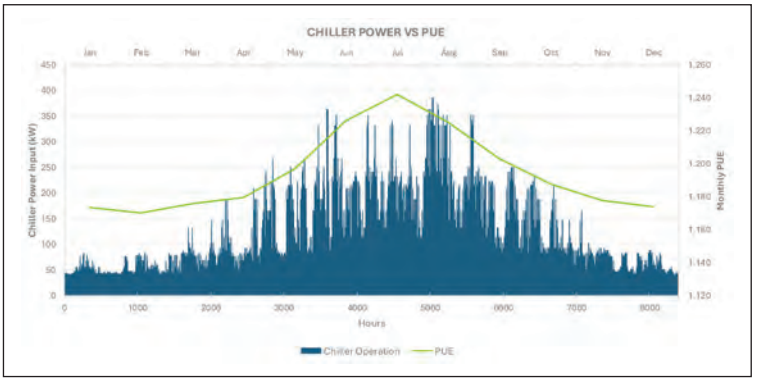

Figure 2: Annual Chiller Power Input vs Monthly PUE

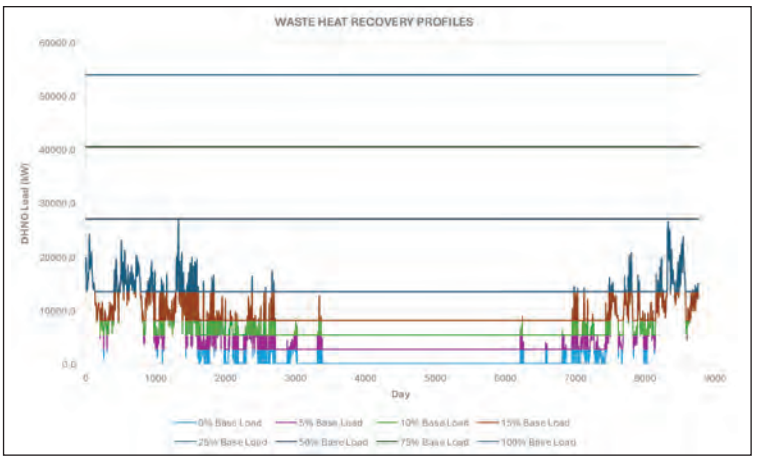

Figure 3: DHNO increased kW base load

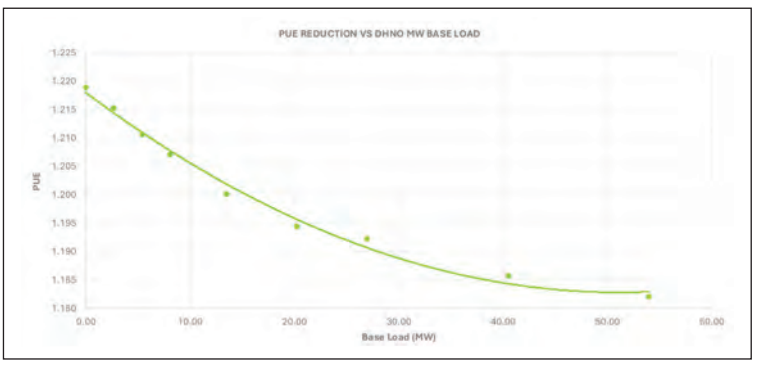

Figure 4: PUE vs % DHNO baseload

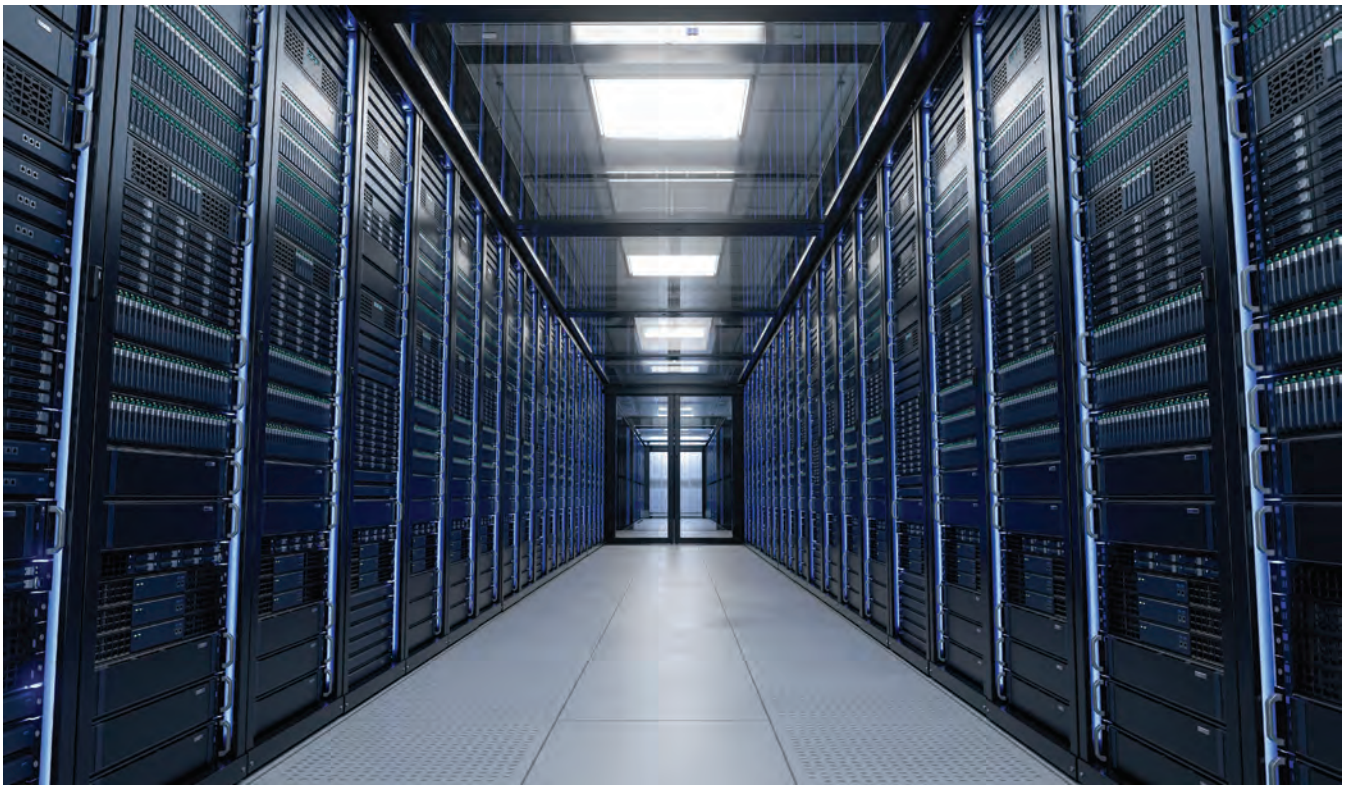

Liquid assets

The significance of liquid cooling for data centre sustainability

BY ALISTAIR BARNES, HEAD OF MECHANICAL ENGINEERING AT COLT DCS

AMONGST the escalating demand for more energy from data centre (DC) customers, maintaining a low Power Usage Efficiency (PUE) rating by DC providers is paramount. With low PUE comes the ability to identify cost saving measures and recognise areas for improvement, whilst also reducing carbon footprints and impacts on the environment.

Yet, as more organisations deploy high performance computing (HPC) and greater IT workloads to support their business requirements, the by-product of heat from power usage is also on the rise. In order for DC providers to ensure that their equipment and systems remain in top condition to serve demand, suitable cooling and heat management is a must.

One effective strategy is building DCs in cold climates and leverage free cooling methods. When the external temperature of a DC is considerably lower than the internal temperature of a facility (as little at 2°C), natural air can be used to directly cool the site and therefore, reduce – if not completely eliminate – the need of mechanical cooling processes. As a result, energy consumption is decreased significantly, and equipment life can be extended reducing embedded carbon. However, geographical constraints such as proximity to key markets and infrastructure mean that not every organisation can benefit from this advantage. For those unable to build in cooler regions, advanced cooling techniques offer a promising alternative.

Among these, liquid cooling has emerged as a particularly effective solution, offering significant advantages over traditional air-based methods. Let's delve into the options and benefits of liquid cooling in modern DCs.

What is liquid cooling?

Rather than utilising traditional air-cooling technology, liquid cooling is a method that uses liquids such as water to conduct heat away from IT equipment.

Compared to air, liquid has a much higher heat capacity (1.004 vs 4.18). Therefore, by using liquid cooling methods, the absorption and transfer rate of heat released from IT equipment will be much higher and more effective when dealing with increasing levels of power.

The different types of liquid cooling

Several liquid cooling methods are available for DC providers. Rack-based liquid cooling circulates coolant through racks to absorb and remove heat directly from servers. Direct-to-chip cooling, including liquid-to-chip methods, targets heat sources by circulating cool liquid through a 'cold plate' in direct contact with components like GPUs. In contrast, immersion cooling submerges entire servers in a thermally conductive liquid, efficiently dissipating heat across all components. Let's delve into this in more detail.

Rack-based liquid cooling:

Rack-based liquid cooling is a method where coolant is circulated through the racks that house servers, absorbing and removing heat directly from the equipment. This approach is highly efficient, allowing for better heat management and enabling higher server densities within DCs. By reducing reliance on traditional air-based cooling, it can significantly lower energy consumption and operating costs.

However, implementing rack-based liquid cooling requires significant upfront investment and infrastructure modifications. Additionally, managing potential leaks and maintaining the cooling system can be complex, posing challenges for DCs not initially designed with this technology in mind.

Immersion Cooling:

Immersion cooling submerges servers in a thermally conductive, dielectric fluid that absorbs excess heat and stabilises IT equipment temperatures.

The process can be single-phase, where the liquid remains in a constant state, or two-phase, where the liquid evaporates and re-condenses to remove heat. Immersion cooling offers significant benefits, such as drastically lowering PUE levels – sometimes as low as 1.10 – by eliminating the need for traditional air cooling systems like CRAC units, which also reduces noise and saves space. However, implementing immersion cooling requires specialized infrastructure and careful management of the dielectric fluid.

Direct-to-Chip Cooling:

Direct-to-chip cooling involves circulating cool liquid through a system that directly contacts the chips and other heat-generating components. This method is more efficient than immersion cooling in dissipating heat at the chip level because the circulating fluid is typically cooler.

Direct-to-chip cooling allows for higher compute densities without needing additional space, making it an excellent option for upgrading existing air-cooled DCs. However, it still requires supplementary cooling for other IT equipment, typically using chilled air, which adds some complexity to the overall cooling strategy.

Hybrid approaches for a balanced PUE

While liquid cooling is an innovative solution, this technology is not yet in a position to completely replace air cooling in DCs. Even if equipment is cooled by liquid, heat will be transferred to it and some of this will be dissipated into a room or surrounding space where air will be required to remove this.

Therefore, a hybrid approach is the best option where liquid and air techniques are used together to offer the best balance of PUE performance.

Furthermore, it is important for organisations to know what options are available to them by working with a DC partner that fully understands their needs. DC providers must be able to review and manage their operations in real-time to ensure efficiency and to meet corporate sustainability goals.

By using modern data analytics tools, businesses can monitor power usage, internal and external temperatures, and electricity usage for cooling to optimise processing loads for cost-effectiveness and proactively monitor equipment maintenance.

Hybrid solutions are increasingly being adopted to combine the strengths of different cooling methods, improve PUE, and pave the way for a more energy-efficient future.

Legacy tech in modern data centres

Navigating the risks of outdated hardware.

BY JAD JEBARA, PRESIDENT & CEO, HYPERVIEW

IN OUR INCREASINGLY digital world, data centres have become the backbone of modern society. The explosive growth of this industry, driven by cloud computing, big data, and artificial intelligence, is reshaping our technological landscape. The scale of this expansion is staggering: the global data centre market size reached $187.35 billion in 2020 and is projected to grow to $517.17 billion by 2030.

While this growth is essential for our digital future, it brings with it a host of critical challenges and hidden dangers that demand immediate attention. Top among these is the often-overlooked issue of aging infrastructure, which sets the stage for potentially catastrophic failures.

The age problem: A multifaceted challenge

The reality of the data centre industry is stark and concerning. According to IDC, the average data centre is nine years old. Gartner highlights this concern further, stating that any site more than seven years old is considered obsolete by current standards. A closer look at the industry reveals that one-third of data centres have at least some proportion of their facilities aged between six and ten years old. Adding this to alarm, approximately 17% of operated stock is ten years old or more. This aging infrastructure is not just an optimal performance issue; it's a ticking time bomb that compromises efficiency, reliability, and safety in the following areas:

- **Inefficient power utilisation:** Older equipment often requires more power to operate, driving up energy costs and potentially overloading electrical circuits. This increased power consumption not only impacts the data centre's carbon footprint but also puts additional strain on already aging electrical systems.
- **Diminished performance:** The deterioration of IT infrastructure over time leads to slower processing times, reduced storage capacity, and increased latency. These performance issues have far-reaching consequences, affecting employee productivity, customer satisfaction, and the overall ability of businesses to meet the demands of modern operations.
- **Increased risk of failures and downtime:** Perhaps the most critical issue with aging infrastructure is the heightened risk of equipment

AFL

“The emergence of generative AI, with its exceptionally large models and truly extraordinary computing requirements, has driven greater and more rapid change in data center networking in the past two years than I have seen in the previous decade.”

Alan Keizer
Senior Technology Advisor, AFL

www.aflglobal.com

failures. This can lead to unplanned outages and downtime, resulting in lost productivity, missed deadlines, and direct revenue losses. Moreover, frequent outages can cause lasting damage to a company's reputation and erode customer trust.

- **Safety hazards:** Aging infrastructure sets the stage for one of the most dangerous potential outcomes: arc flash incidents.

A shocking danger: Arc flash incidents

Among the most serious threats facing modern data centres with aging infrastructure is the increased risk of arc flash incidents. These explosive electrical failures can cause devastating damage to equipment and pose severe risks to human safety.

The U.S. Bureau of Labor Statistics reports that across industries, there are 5-10 arc flash incidents daily, with OSHA estimating that 2,000 workers are admitted to burn centres annually due to these events.

The risk of arc flash incidents in aging data centres stems from several factors:

- **Deteriorating insulation:** Over time, the insulation on electrical components can degrade, increasing the risk of electrical arcing.
- **Outdated safety mechanisms:** Older data centres may lack modern safety features designed to prevent or mitigate arc flash incidents.
- **Increased electrical load:** As data centres expand and add more equipment, the electrical load on aging systems increases, potentially pushing them beyond their safe operating limits.

DCIM: A Transformative Solution

Data Centre Infrastructure Management (DCIM) technology emerges as a crucial solution to these challenges. DCIM is redefining data centre management by integrating cutting-edge technologies to streamline operations, predict and prevent electrical failures, and address industry-wide challenges.

Key benefits of DCIM include:

- **Predictive maintenance:** Leveraging AI and machine learning to predict potential equipment failures before they occur, significantly reducing the risk of arc flash incidents and other catastrophic failures.
- **Real-time monitoring:** IoT sensors provide continuous data on equipment performance, power usage, and environmental conditions, allowing operators to identify and address issues promptly.
- **Optimised resource allocation:** Comprehensive insights enable data-driven decisions about resource allocation, ensuring efficient distribution of power and cooling across the facility.
- **Extended equipment lifespan:** DCIM often reveals that existing equipment can be optimised or repaired rather than replaced, reducing costs and environmental impact.
- **Enhanced safety protocols:** Detailed insights into equipment status and potential risks enable more effective safety measures.
- **Talent shortage mitigation:** Automation of routine tasks and provision of insights help less experienced staff make informed decisions.
- **Scalability for edge computing:** Effective management of distributed infrastructure for evolving technological landscapes.

The path Forward

As the data centre industry continues its rapid growth, addressing the challenges posed by aging infrastructure is crucial. Modern DCIM technology offers a comprehensive solution, providing the visibility and control necessary to ensure both the safety and efficiency of data centre operations.

By embracing DCIM, data centre operators can effectively mitigate risks associated with aging equipment, optimise resource utilisation, improve overall performance and reliability, and enhance safety measures. This not only protects personnel and assets but also ensures that data centres can meet the ever-increasing demands of our digital world.

In an era where our digital infrastructure is more critical than ever, investing in modern DCIM solutions is not just a matter of operational efficiency– it's a fundamental requirement for building a resilient and sustainable digital future.

As we look ahead, the integration of DCIM technology will play a pivotal role in shaping the next generation of data centres, empowering operators to navigate the challenges of rapid growth while maintaining the highest standards of safety, efficiency, and reliability.

APC™

40Years

Make greater

IMPACT

with an industry-leading network

Celebrate 40 years with APC.

Start leveraging four decades of uninterrupted protection, connectivity, and unparalleled reliability with the APC UPS family, a legacy marked by pioneering UPS technology and an unwavering commitment to innovation.

https://www.apc.com/uk/en/product-category/88972-uninterruptible-power-supply-ups/

Life Is On Schneider Electric

Base 16 Technology in the data centre: what are the benefits?

As the data centre market transitions to higher data speeds, it is important to consider what solutions provide the speed and bandwidth necessary to enable those applications that require faster and larger datasets.

BY MICHAEL AKINLA, PANDUIT

HYPERSCALE AND EDGE data centers are installing or upgrading infrastructure fabrics to 400G and even 800G (2x400G) for switch to server and leaf-spine connectivity, for which there are several MMF and SMF transceiver options to choose from. Among these options, many customers are evaluating multimode parallel optics for short-reach network connectivity. 400GBASE-SR8 transceivers over 16 parallel MMF strands (50Gbps on each strand of MMF), require Base 16 MMF structured cabling systems, providing significant value for short-reach connectivity systems interconnecting MOR/EOR switches to servers.

Newer technologies and applications that require faster interconnect, increasing bandwidth, and network capacity require innovative solutions to enable operators to offer leading-edge platforms. There is keen interest in 400GBASE-SR8 systems utilizing multimode transceivers, that for switch-to-server connections provide lower cost and power consumption than SMF alternatives such as 400GBASE-DR4.

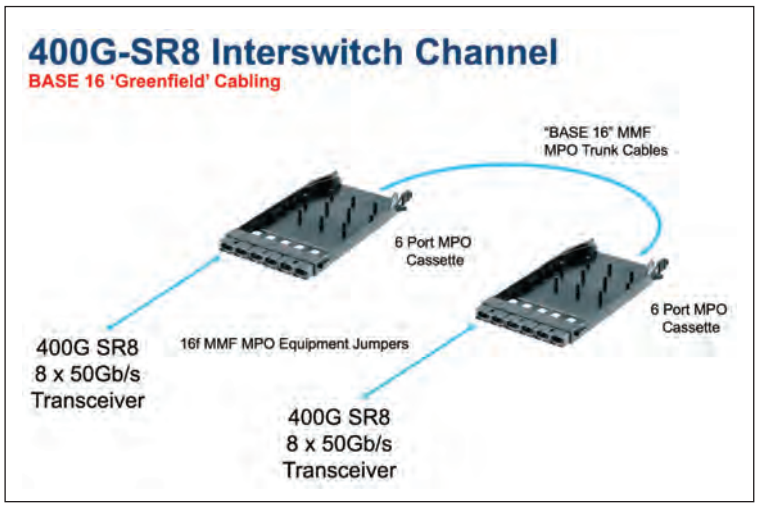

➤ Figure 1. 16 fibre MPO Solution

Base-16 structured cabling, with small form factors cassettes, MPO-16, and Base 16 cabling simplify network deployment in greenfield networks as shown in Fig. 1. To take full advantage of multimode short reach variant, Base 16 interconnection with SR8 transceiver, can be utilised over existing Base-12 structured cabling infrastructure conversion cassettes.

We are currently seeing single mode traffic in terms of some customers moving from 400G to 800G (2x 400G), however, our customer interactions indicate that the multimode 16 fibre MPO modular cassette solution certainly has a place in the market, particularly for MOR/EOR switch to server interconnections due to the previously mentioned advantages. This is especially true in respect of preterm assembly offering speed of deployment, with highly flexible cassette configurations.

Expanding capacity

The key benefit that enabling high density breakouts of 50G server ports using SR8 is the level of expansion if offers. When you have Base-16 fibre from an SR8 perspective it offers 50G breakout, for example, from a 400G SR8 optic, being delivered as a 32 port 1RU modular 400G switch, will then enable 256, 50G breakout ports in terms of a high Radix.

A benefit of this approach is a reduction in the number of optical links and supporting cable plant (patch panels and connectors) as the deployment is 400G to 400G and then using breakout technology to get the 50G at the top of the rack (TOR) server end without using an actual TOR switch. This allows operators to deploy less fibre infrastructure into the data center environment. A further benefit therefore is the fourfold increase in respect of aggregator switch bandwidth, and in some situations a lower initial installed cost, depending on the SKU.

COOLING | HEATING | VENTILATION | CONTROLS

Ask ME about our IT Cooling solutions for Data Centres

Multi Density

Scan here

Data centres are now a crucial part of modern life and Mitsubishi Electric's **Multi Density** system makes IT Cooling easier than ever. Combining the efficiency of VRF with high-performance close-coupled air conditioning, Multi Density offers precise temperature and humidity control for high-density applications. It ensures optimal conditions for your IT equipment whilst helping increase energy efficiency and lowering emissions.

With features like advanced inverter technology, a small footprint, and pipe runs up to 165m, Multi Density is perfect for cooling high-density racks and blade servers.

Discover more about our **Data Centre Solutions** at: **les.mitsubishielectric.co.uk**

MITSUBISHI ELECTRIC

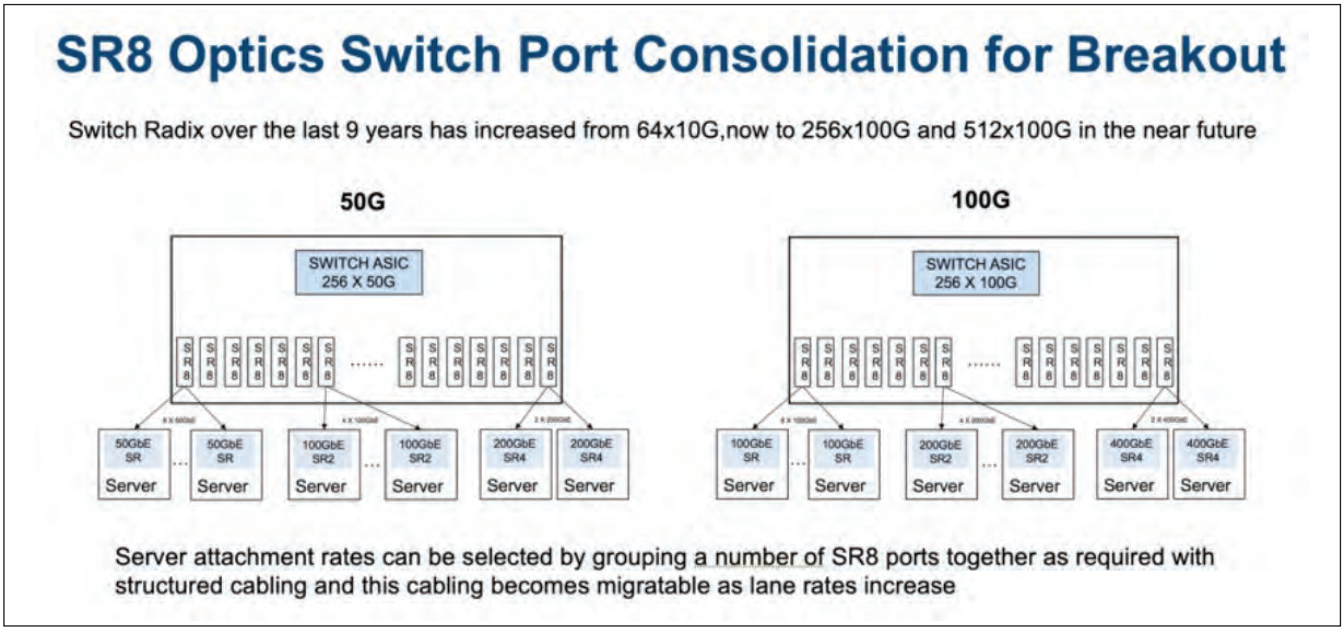

➤ Figure 2. SR8 Optics Switch Port Consolidation for Breakout

Moreover, the replacement of several TOR switches with a few MOR/EOR with more capable and efficient switches that produce significant reduction in power consumption.

The 400G breakout to 50G servers appears to be a sweet spot from a multimode perspective, using a variety of options such as the conversion cassettes, for customers who want to use existing infrastructure across legacy Base-8 and Base-12 fibre cable plant in conjunction with SR8 Base-16 optics.

400G DR4 transceivers can also breakout into 4 streams of 100G. This can be useful to implement the spine-and-leaf connections between switches. Breakout cassettes as the one shown in Fig. 3 can be used. However, contingent on the deployment length and application, SR8 multimode can provide advantages compared to single mode fibre DR4.

SR8 technology helps to simplify the server pod build providing different options. Data centers can deploy patch panels above the servers which breakout of the 400G transceiver with Base 16 MPO coming into the fibre panel into duplex LC ports, and then use patch cable to connect to the servers at 50G as shown in Fig. 3. Another option is to breaks out, with a hydra, from the Base 16 MPO delivering 50G into the server environment, together with an EOR/MOR Switch with 400G ports that break out onto the patch panel and feeds the servers. These deployments enable pre-terminated overhead

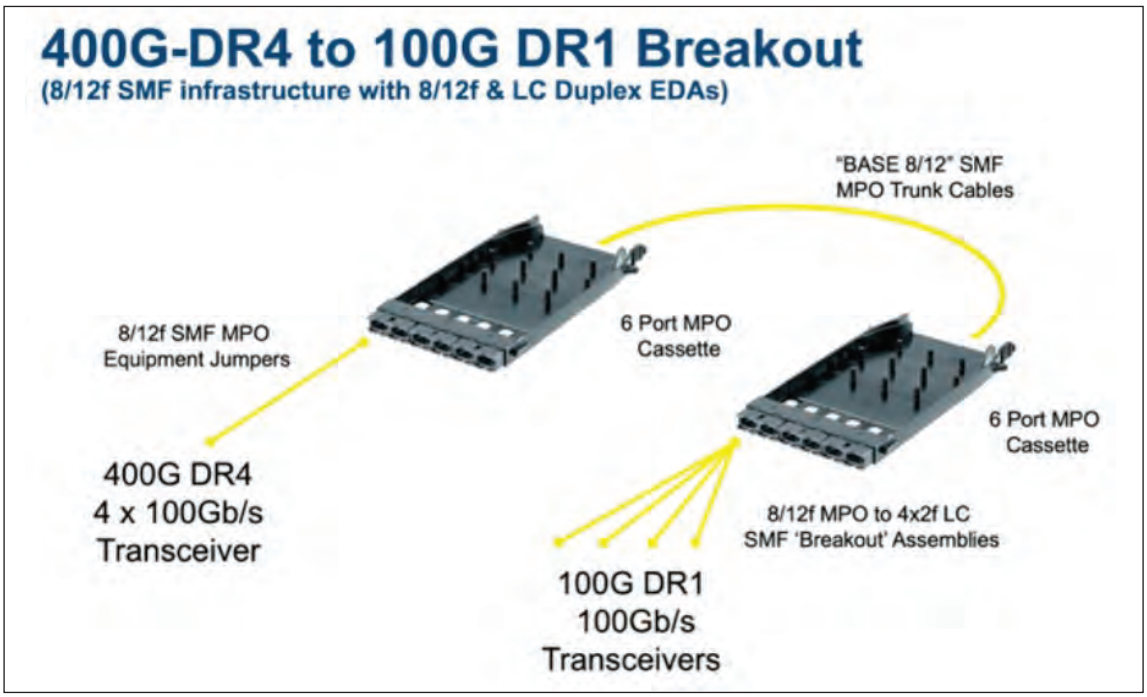

➤ Figure 3. 400G DR4 to 100G DR1 Breakout

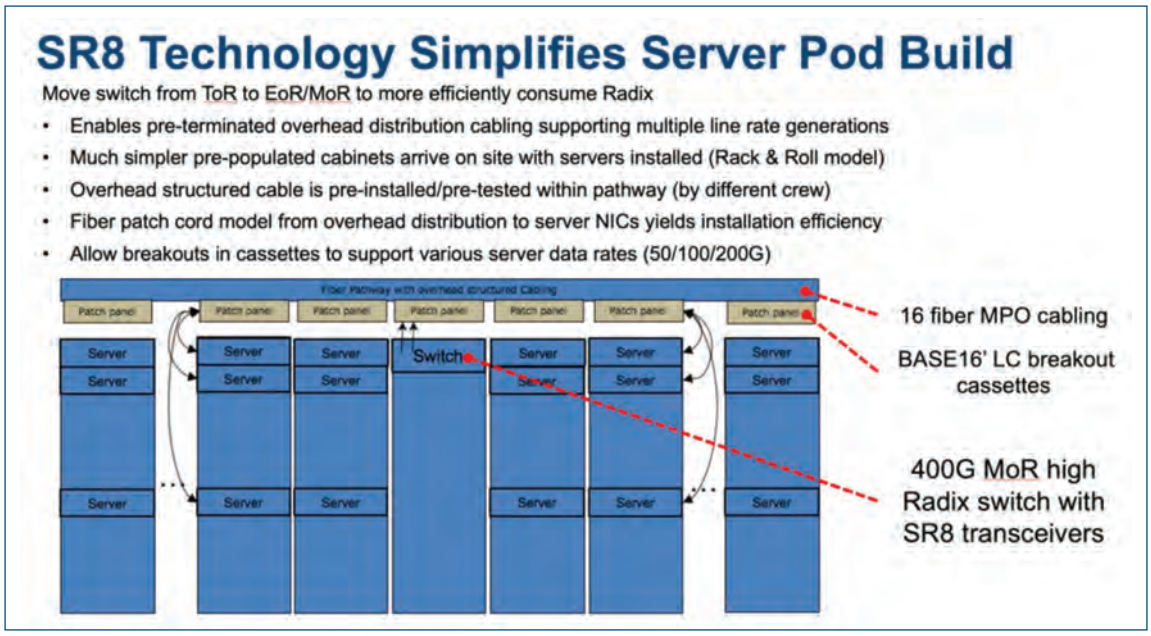

➤ Figure 4. Simplify Server Pod Build , SR8 Technology Simplifies Server Pod Build

cable providing simplified cable distribution in pre-populated cabinets. Overhead infrastructure is advantageous in a rolling rollout environment allowing easier access to end of life cabinets and hardware (Fig. 4).

Conclusion

As the data center market transitions to higher data speeds, we need to consider what solutions provide the speed and bandwidth necessary to enable those applications that require faster and larger datasets. Essential to this is the consideration of how the devices that enable this connectivity to interoperate with existing networks, and how their introduction will move the capabilities forward. Base 16 has a place in that discussion as business drivers for data center operator include new network upgrades and increasing bandwidth while reducing latency.

New 200kW Data Centre Load Bank ideal for Integrated Systems Testing

Crestchic has announced the launch of a new 200kW loadbank, which has been built specifically for use in data centres.

WITH MORE THAN 40 years of experience in the design, manufacture, sales and rental of load banks and power testing solutions, Crestchic have extensive experience in the data centre sector. Following comprehensive customer research, the team has launched a stackable 200kW loadbank, designed to make heat load testing of server halls and electrical infrastructure simpler, faster, and more efficient.

Paul Brickman, Commercial Director at Crestchic, explains, “Having served the sector for many years, we have been able to collaborate with key customers to get to the crux of what they require from a load bank. By taking that experience and insight back to our engineering team, we have been able to build a solution that fits an exacting specification. The result is a machine that provides an ideal solution for integrated system tests, is easy to operate, and easy to store. As higher data centre density becomes the norm, testing cooling systems and performing integrated systems tests is increasingly important - this product meets that need.”

At just 1159mm x 876mm x 983mm the load bank has a small enough footprint for operators to link multiple machines and strategically place them at intervals around the hall. This makes it possible to closely mimic the heat generated by servers, replicating airflow patterns and rack configuration such as hot and cold isles. Multiple load banks can be configured and controlled using just one LC10 handheld digital controller, enabling fully remote operation and simple management of incremental loads. This allows operators to use up to 20 linked load banks (4MW) to verify that the cooling installation is operating to specification in the server hall and within safe parameters, to check the redundancy of the system should a cooling unit fail, and to verify that there are no hot spots that may require reconfiguration.

In addition to heat load testing, the units are invariably connected to Power Distribution Units (PDUs), facilitating testing of the data centre’s electrical infrastructure, including power to servers, switches, bus track power units, feed tracks, tap offs, and other devices. With heavy duty casters, the new 200kW data centre loadbank has been built for easy manoeuvrability and can be positioned at each PDU in the server hall, for full system testing.

The units have one of the lowest Delta T rises on the market at below 100°C (at the exhaust face), helping to ensure a more accurate simulation of real-world conditions, and providing a more accurate test of the cooling system’s ability to handle the actual heat load it will encounter during operation. The load bank boasts a compact, stackable design, enabling the units to be safely stacked two high, reducing the storage footprint when stored. They have also been built to ensure that there are no awkward shapes or protrusions, which could be subject to knocks and bumps. Noise levels are lowest in class at 83dBA at 1m, an important consideration when operating inside

Paul adds, “With the sector now expecting 99.999% uptime as standard, it is important that data centre operators test their critical systems. Using these units for integrated systems testing pushes the systems to their limits by simulating real-world scenarios like full load operation, power outages, and cooling system failures. As well as helping to identify any weaknesses or bottlenecks before critical IT equipment is deployed, the test generates valuable data on system performance under different loads, providing a baseline for future reference that can support maintenance and troubleshooting down the line.”

Crestchic offers a range of solutions for testing backup diesel generators and UPS; heating, ventilation, and cooling (HVAC) systems; electrical bus and distribution cabling, and PDUs. Load testing at commissioning stage, and at regular maintenance intervals, can help to ensure that systems remain operational and reduce the risk of unwanted and expensive downtime.

200kW units are available for short and long-term rental. They are also available for immediate purchase, with delivery by January 2025.

For more information, visit: **crestchic.com**

wesco | anixter

CHATSWORTH PRODUCTS

Protect

Global Expertise in Protective Data Centre Solutions

Delivering value at every stage of data centre development and operations, Wesco Anixter brings reliable expertise alongside a global ecosystem of strategic partners. Together with CPI, we provide advanced rack and cabinet solutions to protect your cabling infrastructure, ensuring robust security and reliable connectivity to keep your network running efficiently.

Scan to learn more about CPI's rack and cabinet solutions.

We build, connect, power and protect the world.

Wesco.com/datacenters-emea

Trusted manufacturer of essential components for data centres

Essentra Components, a global market leader with over 65 years of expertise, has become a trusted manufacturer of high-quality components for many industries worldwide, including data centres

FROM ACCESS HARDWARE to cable management, Essentra Components delivers a comprehensive range of solutions that help streamline data centre operations, ensuring they meet the highest standards in durability, security, and functionality.

Why Choose Essentra Components for Data Centres?

Essentra Components' mission is simple: to make it easier for customers by providing top-tier products and hassle-free service. With operations in 28 countries and an expansive portfolio of 45,000 products, we ensure fast, reliable access to the components data centres need. Over 1 billion parts are kept in stock, ready for immediate dispatch, ensuring that our customers get what they need when they need it. For you, that means our manufacturing capabilities, experience, customer focus, and global infrastructure allows us to deliver best in class solutions, tailored to your needs

Our heritage

Essentra Components began in 1955 in Kidlington, Oxfordshire, and has since grown into a global leader in plastic injection moulding, reaching 28 countries and spanning four continents. Rebranded in 2013, Essentra demonstrates quality and expertise across various industries.

The company's growth has been fuelled by strategic acquisitions, expanding its product range and global presence. Acquisitions include BMP TAPPI, Wixroyd, Hengzhu, Mesan, Skiffy, Richco, among others, enabling significant investments in production and service facilities worldwide.

Here's a closer look at how our core components support data centre operations.

Essentra Components offers a vast selection of essential components for the demanding needs of data centres. From cable management and locking solutions to PCB hardware, our range ensures reliable outcomes at every stage of your data centre lifecycle, from inception to completion.

Our range of access hardware is designed to optimise security in data centres, where sensitive information must be protected against unauthorised access. Our Electronic Locking Systems offer advanced monitoring and tracking features, ensuring that access to server cabinets and other equipment is closely controlled. These systems are ideal for data centres, as they allow for real-time

DCS AWARDS 2025

CELEBRATING 15 YEARS OF SUCCESS

The DCS Awards: 38 Categories across 5 Themes

22 MAY
Leonardo Royal Hotel
London St Pauls

KEY DATES:

14 FEBRUARY: NOMINATIONS CLOSE

21 MARCH: SHORTLIST ANNOUNCED

24 MARCH: VOTING OPENS

23 APRIL: VOTING CLOSES

22 MAY: AWARDS CEREMONY

NOMINATE NOW!

SPONSORSHIP PACKAGES

The DCS Awards offer extensive branding and sponsorship opportunities through online advertising in our Datacentre Solutions & Digitalisation World publications, and of course at the awards ceremony itself.

For sponsorship opportunities and/or to book your awards table please contact: awards@dcsawards.com or call +44 (0)2476 718970

NOMINATE: https://dcsawards.com/nominate

Supported by

The data centre trade association

monitoring and control, adding an extra layer of security.

Featured products for data centres include:

- Cabinet locking solutions; including swing handle locks, rotary latch systems, swing handle locks
- Hinges for doors and panels
- Feet and castors
- Sealing gaskets and edge protection
- Handles
- Cable entry
- Panel fasteners
- Cable management
- PCB hardware

Supporting Innovation in Data Centres

Essentra Components understands that the needs of data centres are constantly evolving. That's why we offer prototyping and custom manufacturing solutions to accommodate unique project requirements. If you're developing a new cabinet design or need specific hardware components, we can assist with custom-built solutions, prototyping, and even 3D printing. This service helps data centre operators and designers innovate without being limited by standard components, allowing them to create tailored solutions that fit their specific applications.

To further support the design and prototyping process, we provide free samples and CAD models of our standard parts, so designers can visualise how each component will fit into their project.

Why Essentra Components?

At Essentra Components, we bring over 65 years of expertise to every project. With a team of 2,500 experts across the globe, we provide solutions to industries ranging from data centres to automotive, medical devices, renewable energy, and more. Our dedication to quality and innovation ensures that our components meet the highest standards, supporting emerging industries and established markets alike.

Supporting your sustainability goals

We're making it easier for you to achieve your sustainability goals by developing products that reduce your impact on the environment.

How?

- Developing sustainable component solutions by using more sustainable materials and making our products recyclable
- Circular Plastics Alliance commitment to use at least 20% recycled content by 2025
- Introduced over 50% recycled content in many of our LDPE & PP ranges
- Assisting you by providing information to support your scope 3 reporting

Our commitment to you and our plant

- Committed to setting science-based targets for emissions reduction, targeting net zero in our direct operations by 2040, and in our value chain by 2050
- Targeting all of our sites to achieve zero waste to landfill by 2030 at the latest
- Supporting a circular economy by ensuring 100% of our packaging is reusable, recyclable or compostable by 2030
- Aiming for 50% recycled content in our packaging materials by 2030

Final thoughts

In a world where data is invaluable, data centres must be equipped with components that are durable, secure, and designed for efficiency. Essentra Components stands as a reliable partner, dedicated to delivering quality solutions and ensuring your data centre operations run smoothly. With our vast selection of products and dedication to customer satisfaction, we are ready to meet the evolving demands of the data centre industry.

Whether you're expanding an existing facility or building a new one, Essentra Components has the products and expertise to support your project from start to finish.

For further information or to discuss custom project requirements, our team is ready to assist with expert guidance and tailored recommendations.

Explore our full range and discover helpful guides, expert tips, and more at www.essentracomponents.com

Visit our Knowledge Centre for insights on choosing the perfect products to fit your needs, and feel free to reach out – our expert team is here to help.

For any questions, contact us at sales@essentracomponents.co.uk or call 44 (0)345 528 0474.

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.

The DCA Update

By Steve Hone, CEO The DCA

THIS ISSUE of Data Centre Solutions has a number of themes – Hybrid World, Connectivity & Networking and Modular DC's.

DCA Partners CMP Products, Harting and Delta have contributed three informative pieces for The DCA Update. I'd like to personally thank Ajay Kareer - Data Centre Market Manager, HARTING Ltd, Arno de Ruyter - Marketing Manager, Telecom Power Solutions ICTBG EM / DELTA and Lee Frizzell - Technical Director, CMP Products for taking the time to produce these articles.

The DCA's – Data Centre Transformation Conference 2024
Thank you to all that attended the conference in Birmingham on the 22 October.

It was a 'sell out' with no empty seats in the lecture theatre!

We have set the date for 2025. 21 October 2025 – so hold that date in your diaries. You can find out more here: https://dctransformation.co.uk/

Come and meet The DCA in November and December 2024
Data Centres Ireland 2024
20 November 2024, Dublin
DCA representatives will be attending the event, representing DCA Members & Partners and speaking on panels
Stand No 143

PowerEx Live London
12 December 2024, central London
The DCA have stand and will be representing DCA Members & Partners and in DC Sector along with participating in the conference programme.

If you would like to arrange a meeting at either event please drop me a message – Steveh@dca-global.org
Find out more about The DCA or email us mss@dca-global.org

Best regards
Steve

Modular telecom data centers:

A win-win in a rapidly evolving market

BY ARNO DE RUYTER - MARKETING MANAGER, DELTA ELECTRONICS

THE TELECOM SECTOR is rapidly evolving, driven by the convergence of telecom and IT networks, IoT, and the expansion of 5G and cloud services. Traditional brick-and-mortar data centers (DCs) can't keep up with these demands. Delta Electronics offers a flexible, scalable, and cost-effective solution with our prefabricated modular data centers (MDCs), manufactured in Croatia. These MDCs enable rapid deployment in diverse locations, from urban rooftops to remote sites. MDCs support next-gen network expansion by reducing latency, improving energy efficiency, and allowing easy adaptation to fluctuating demand.

Case studies with Eurofiber and Bouygues Telecom highlight MDCs' critical role in optimizing network performance, reducing time-to-market, and enhancing sustainability. As challenges persist, the modular approach offers the agility telecom operators need to remain competitive and sustainable in a fast-changing environment.

CASE STUDY: **Eurofiber - MDCs support flexibility for a fiber-optic network operator**
Eurofiber is a leading operator of open digital infrastructure in Benelux and the North of France. Headquartered in the Netherlands, the company's fiber network is more than 70,500 kilometers long and is extended by 50 kilometers every week. To maintain optimum connection speeds throughout its network, even as it expands, Eurofiber uses inline amplifiers (ILAs) to boost signal strength. An ILA is typically placed every 75 kilometers. Among the technical solutions used in conjunction with the ILAs are Delta's Xubus Nodes.

Xubus Node is a modular data center (MDC) characterized by high performance and a high grade of mobility, which can be deployed quickly and easily in standardized containers of various sizes. This includes sites where the supporting infrastructure is not fully in place, but computing power is required. Xubus Node comes in five different configurations based on the required IT workload. In addition to fiber networks, they are also frequently used in the expansion of 5G networks. For the expansion of Eurofiber's fiber optic "highway" between Amsterdam and Brussels,10 Delta rolled out a custom-designed version of its Xubus Node, optimized to provide the lowest possible operating costs.

In addition to streamlined UPS modules and a highly energy-efficient ceiling cooling system, Delta integrated a rooftop solar system to enhance energy efficiency.

A smart access control system for doors and racks controlled via mobile phone created further efficiencies. It means Eurofiber can manage access to its many DC sites remotely, without the need for on-site staff.

Delta's DCIM (Data Center Information Management) system allows Eurofiber to monitor their data centers (DCs) centrally in real-time, including their power consumption and fault and fire monitoring. The turnkey solution Delta designed and rolled out provides Eurofiber with a tailor-made ILA DC that combines high performance with great stability, reliability and flexibility as its network continues to expand at speed.

CASE STUDY: Bouygues Telecom fiber network expansion

Before selling its fiber network assets in France to telecom infrastructure specialist Cellnex,19 French multi-service operator Bouygues Telecom worked with Delta to expand its edge network infrastructure for 5G deployment.

At one of the first sites, Châteaurenard near Avignon in Southern France, Delta helped Bouygues set up a prefabricated modular data center (MDC) as a model to be replicated across the whole of France. Choosing a modular design meant it could be completed a year faster than a traditional brick-and-mortar data center (DC).

The MDC units were manufactured in Croatia, at Delta's center of expertise for modular DCs. The on-site team also pre-tested and pre-assembled most of the systems, minimizing efforts on-site in France. Thanks to the modules being pre-integrated, Delta was able to work with local contractors in France to deploy the DC, rather than bringing in an installation team from Croatia.

Taking the prefabricated approach also reduced the number of suppliers Bouygues needed to work with, as very little physical construction work had to be undertaken. Along with reducing the complexity of managing the project, this approach also led to quality improvements.

Conclusion

Technology is only one of the challenges facing the telecom sector. With more market consolidation expected and competitive pressures unlikely to abate, making flexible infrastructure investments and decisions will continue to be vital for telecom network operators. The high level of agility that can be achieved with modular approaches, such as those offered by prefabricated MDCs, enables operators to respond more effectively to a fast-paced market and technology environment.

A staggering 181 zettabytes of data will be created, captured, copied and consumed worldwide in 2025 — almost twice as much as in 2022.1 The major drivers for these skyrocketing volumes are the proliferation of internet-of-things (IoT) services and the expansion of mobile connectivity and cloud services, culminating in the rollout of 5G mobile networks.2 These — along with future 6G network technology3 — will enable some of the most sophisticated applications we have ever seen, especially as high band services in the millimeter wave spectrum come online. Providing connection speeds of 3Gbps and more, this will enable low-latency services such as self-driving cars, high-frequency trading, virtual reality and real-time streaming.

Read more here: https://dca-global.org/file/view/15249/modular-telecom-data-centers

Securing electrical safety

LEE FRIZZELL, TECHNICAL DIRECTOR OF CMP PRODUCTS talks to Data Centre Solutions about why it's still of vital importance that the traditional elements of electrical safety and support systems are given as much consideration as the millions spent on cyber-security and online systems in ensuring system integrity in data centre installations.

CABLE CLEATS play a vital role in any electrical installation – both in terms of ensuring the continuation of electrical supply and protecting high value systems during a short-circuit scenario.

"When you consider the millions spent creating and maintaining networks and electrical systems in data centres, it would be a fair assumption to say that currently, cable cleats wouldn't feature particularly high on a list of products included amongst the vital elements of the build.

However, the importance of cable cleats must not be underestimated. They are safety critical products that need to be correctly specified in order to safely secure electrical cables not only during normal operating conditions, but also ensuring they are able to withstand extreme dynamic electromagnetic forces during a short circuit without sheath damage.

Effective cleating of cables, using a tested, proven product ensures the continuation of potentially vital electrical supplies and prevents cables leaping out of their containment systems during a short circuit, which would pose a very real threat to the system integrity and the lives of those in the vicinity of the affected installation.

Put simply, for any electrical installation to be deemed safe, cables need to be protected and restrained by devices proven to withstand the forces the cables generate - especially during a short circuit - and this is exactly what cable cleats are designed and tested to do. Without effective cable restraint, there's risk of damage to expensive cables and cable management systems, and the previously mentioned threat that poorly restrained live cables pose."

So what are the key things to consider when it comes to specifying cable cleats?

A question of trust

There are lots of cable cleats out there and it's fair to say they can look very similar and claim to have very similar attributes. Therefore, it's important to specify with confidence. You'd think twice about buying an expensive new phone from an unknown telecoms manufacturer, so don't do the same with cable cleats. Look at the company behind the product; their technical expertise; the post-sales support they offer; where else have their cleats been used

service life – so be certain you're making the correct specification decision for the right reasons.

Testing, testing

Will the cable cleats you're considering be up to the job?

Some manufacturers will highlight the strength of their products based solely on a static mechanical tensile test. Unfortunately, this can be misleading as the forces applied in this type of test are applied in a slow and controlled manner, unlike in a short-circuit fault, where dynamic forces are applied almost instantaneously and oscillate in every direction.

A short circuit test really is the only reliable way of knowing a cable cleat is capable of withstanding a specific set of fault conditions and is essential in order for anyone, anywhere to specify with confidence.

Make space

It's not just a question of getting the right cleats – they need to be correctly spaced to ensure they perform properly. Too close and you've over-specified. Too far apart and they won't be strong enough to match the performance outlined in the certified short circuit test.

A trustworthy manufacturer will advise on cleat spacing and demonstrate the force calculations

material corrosion also needs to be considered. Stainless steel is the material of choice for the vast majority of cable cleats and fixings due to its non-magnetic and excellent corrosion resistant properties, with 316L being the most commonly used variant. Its corrosion resistance is the result of its chromium content, which reacts with oxygen and forms a self-healing impervious layer of chromium oxide on the surface of the steel.

Final considerations

The final three things to verify before specifying cable cleats are very much down to common sense:

- **Fire safety:** Make sure the cleats being specified have the same flame retardancy as the cable they're securing.
- **Operating temperatures:** Most cable cleats are designed for use in ambient temperatures ranging from -60°C to +60°C and with cable conductor temperatures up to +90°C. Some installations require cable cleats proven for higher temperatures.
- **UV resistance:** Metallic cable cleats are impervious to UV attack. Composite and polymer products aren't. If they're likely to be exposed, they should be supplied in materials containing

installation, but their performance is paramount to its performance, safety and the longevity. Therefore, specify with certainty and don't be tempted to cut corners or costs."

CMP Products is the leading manufacturer of cable glands, cleats and accessories. Its cable cleats are in demand around the world and have been used extensively in data centres around the world.

To find our more visit www.cmp-products.com

Reliable, energy efficient connectivity solutions for data centres

BY AJAY KAREER, DATA CENTRE MARKET MANAGER - HARTING LTD

THE WORLDWIDE DATA CENTRE market is experiencing explosive year-on-year growth as our reliance on remote working, AI and the Internet of Things increases at a staggering rate. In addition, the changes to our working lives have meant that businesses and individuals need reliable access to data to allow them to embrace flexible ways of working. Therefore, as we become more reliant on remote or hybrid working models, it's essential that data centres run as smoothly and efficiently as possible.

Data centre power outages can happen for various reasons such as weather conditions, network failures, human error and software issues. However, they can also occur due to power infrastructure problems created inside the data centre from either generator, Uninterruptible Power Supply (UPS) or Power Distribution Unit (PDU) failures.

The International Data Corporation reports that energy consumption per server is growing by around 9% per year globally. Despite servers getting more compact to save installation space, their improved performance increases their energy requirement. As a result, energy consumption costs can be more than 50% of the total data centre operating expenses (OPEX). It's therefore essential to invest in and manage each part of the critical infrastructure in the data centre to ensure energy efficiency and reliability. Cable assemblies distribute power from the data centre's UPS to the PDUs. These assemblies consist of a cable between one or two connector hoods. Inside the connector is an insert or multiple inserts where the conductors from the cable are terminated. The connector hoods then mate with a matching housing wired to the PDU and/or UPS.

When a cable assembly is designed and manufactured using automation, human error is massively reduced. If the same connections are handmade or field wired, the chance of error increases, potentially risking catastrophic issues either during the initial power up or during the operation of the data centre. This in turn can result in hours of expensive skilled labour spent troubleshooting as well as the downtime costs of the rack, PDU or entire data hall not functioning. If designers hardwire the conductors inside the cable, a skilled electrician is needed to disconnect and reconnect the hard-wired PDU. Using cable

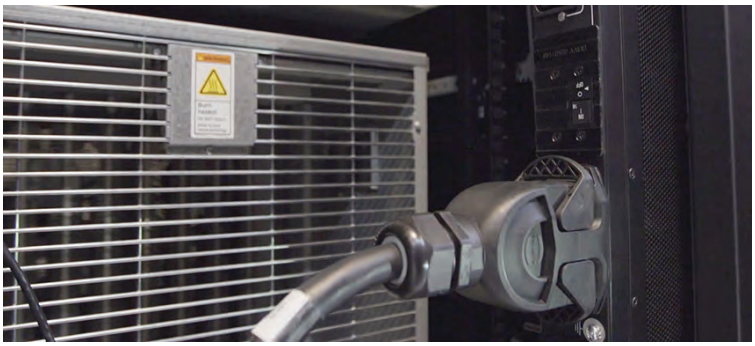

assemblies means there is no need to hire an electrician and, since everything is pre-wired and pre-tested, wiring errors are virtually eliminated. As energy costs can account for more than 50% of the total operating expenses of a data centre, one important ongoing challenge is to improve the energy efficiency of its infrastructure. To calculate the exact effect of power usage from connectors in data centres, HARTING has compared the power consumption of three different connector solutions in its independently accredited test laboratory. One of the connectors tested was the HARTING Han-Eco®. The other two were CEE (IEC 60309) plugs from different manufacturers.

The results showed that the Han-Eco® connector reduced power wastage by up to 50%, compared to the other two brands of IEC connectors, by using low-impedance contacts. These contacts reduce the power lost in connections and significantly improve the Power Usage Efficiency (PUE) of data centres. Depending on the electricity price, which differs regionally and worldwide, different monetary gains can be realised. As an example, one hyperscale data centre with 15,000 racks can achieve annual power consumption savings of around £90,000. These calculations are based on the average EU industrial prices from 2020, so potential savings will be even more dramatic when we consider how much energy prices have increased over the past four years.

The Open Compute Project (OCP) is focused on the redesign of hardware technologies for IT infrastructure. The goal of the working group is to make data centres more efficient, more flexible and more quickly scalable, via an open exchange of ideas, specifications and other intellectual property to maximise innovation and reduce the complexity of technical components.

In a data centre, power shelves provide power to IT equipment. The Rack & Power Project Group within the OCP initiative is focused on standardising racks and making them easier to integrate into the data centre infrastructure. These designs, called the Open Rack, began worldwide installation at the beginning of 2023. As a lead author and initial connector partner in the standardisation process, HARTING has now developed the third version of the Open Rack (ORV3), called the ORV3 OCP Input Power Connector. In line with the OCP's goal of optimising efficiency in the construction and scaling of data centres, the Han® ORV3 enables a more compact design for the entire infrastructure thanks to its shallower rack system. Connectivity technology is constantly being refined and developed, and new Smart Connectivity solutions are designed to improve safety, identify faults, and ensure systems within data centres are working efficiently.

One of the most important additional functions powered by Smart Connectivity is the signalling of the mating state. The mating state can indicate a range of different parameters, including if the connector is electrically connected and whether it is mechanically locked. It can also indicate if the connector is overloaded and monitor whether environmental parameters such as temperature and humidity are within the permitted range. Connectors are currently identified by using electrical contacts as coding pins, with the control system determining which attachment is plugged in. However, this method has its limits, especially with large flexible systems.

The latest solution identifies the connector with the help of a bus system and microcontroller or alternatively via NFC (Near Field Communication). This gives each connector a unique ID which is assigned to the corresponding attachment or tool. As a result, even simple components such as lamps, door contacts or analogue sensors can be identified. To learn more about HARTING and their range of connectors and cabling solutions for data centres, please visit https://www.harting.com/en-GB/data-centres or email: salesuk@harting.com