DCS DATACENTRE SOLUTIONS

DEVELOPING DIGITAL INFRASTRUCTURE IN A HYBRID WORLD

SPECIAL EDITION

HTTPS://DATACENTRE.SOLUTIONS

AI – THE GREAT BALANCING ACT

CHALLENGES OR OPPORTUNITIES FOR THE DATA CENTRE?

WITH MARC MARAZZI – V.P. LEGRAND DATA CENTER SOLUTIONS EUROPE

legrand® | DATA CENTER SOLUTIONS

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

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

WE ARE AI READY, ARE YOU?

Legrand, Your AI Infrastructure Ally

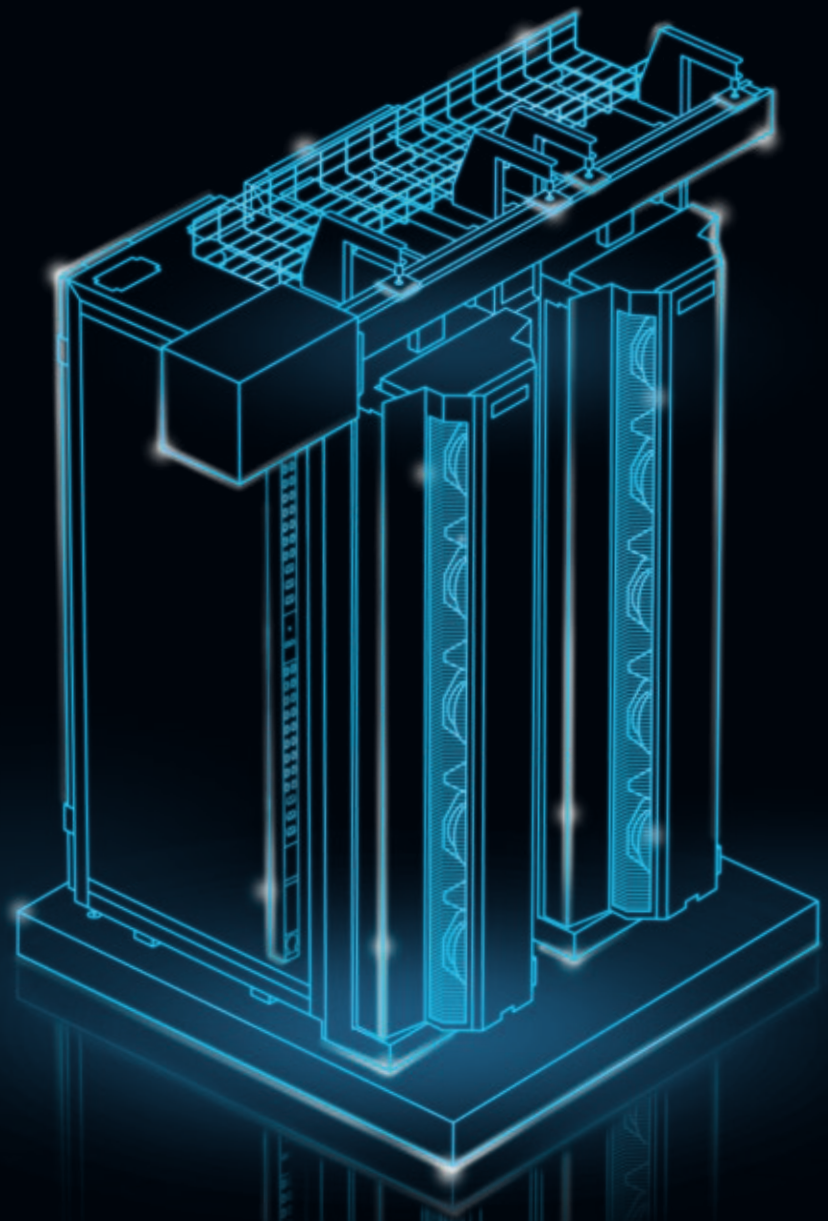

POWERED BY EXPERT BRANDS

MINKELS | Raritan | Server Technology | Starline | USystems

www.legrand.com/datacenter

legrand®

VIEWPOINT

By Phil Alsop, Editor

The best of times, the worst of times?

APOLOGIES to Mr Charles Dickens for borrowing from his opening to 'A Tale of Two Cities', but this phrase seems to encapsulate the current digital revolution that threatens every bit as much disruption as did the French Revolution back in 1789. Indeed, the embryonic AI revolution could just eclipse any previous political turmoil as it offers so many benefits, but so many potentially catastrophic downsides.

At the start of a year when some significant general elections could well be influenced, if not won, by the (mis)use of digital technologies, the data centre industry finds itself with twin challenges. Firstly, it has to respond to this AI explosion, providing a whole new level of data centre density and capacity and, for now at least, power demand. Secondly, it needs to make further, significant progress on its road to Net Zero.

My favourite question to interviewees of our video interviews which feature on the DCS website runs along the lines of: 'Can the data centre industry equally successfully meet the ever-expanding demand for digital infrastructure, with AI yet to come, at the same time as meet its sustainability commitments – whether those of 2030, to which much of the industry has committed or the more 'leisurely' 2050 general global deadline?' Put more sympathetically – should the industry be clobbered for its carbon footprint if all it is doing is providing a service to those who seemingly want to spend their whole lives in a connected and increasingly virtual world?

We can perhaps park this debate for other occasions, but it would nevertheless seem fair to observe that the industry as a whole needs to continue, if not speed up, its sustainability innovation roadmap, if only to ensure that it can deflect the environmental brickbats thrown its way as, irony of irony, eco-warriors tweet their disapproval of 'power-guzzling' data factories!

All of which musings (or ramblings if you like...) brings me to recommend to you the articles in this DCS Special Issue, which demonstrate quite clearly that the industry is no stranger to innovation and does take seriously the responsibility it has to both its digital paymasters and the planet.

The French Revolution didn't turn out too well for royalty or the nobility, but it did see the forging of modern France. It will be fascinating to see how our ongoing Digital Revolution plays out, and who are the winners and losers over time. I don't foresee that, any time soon, data centres will only contain tumbleweed and a few rusty racks and cabinets, but they may well be in the firing line as or when a digital backlash gains significant momentum.

How likely is that? Well, all I will say is that the history books (especially the most recent) are full of surprises. For now at least, let us enjoy the best of both worlds – the digital and the sustainable - and commend the data centre industry for the continuing innovation, as found in these pages, which makes both more than possible into the future.

CONTENTS

ISSUE 0 2024

DCS DATACENTRE SOLUTIONS

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.

The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. ISSN 2756-1143 (Online)

AI – the great balancing act

MARC MARAZZI, VP AT LEGRAND DATA CENTRES EUROPE, talks to DCS about AI and the related HPC challenges facing data centre owners and operators, explaining the ways in which Legrand's infrastructure portfolio is designed to address this rapidly changing data centre landscape, and how the company continues to evolve to ensure an optimised customer experience.

EVERYONE AGREES that the AI explosion, with Generative AI applications already giving us some clue as to the likely impact, will be transformational. There's considerably less agreement when it comes to the how, where and when, in terms of the exact ways in which AI will impact both the office experience and our leisure time. Nevertheless, we all want to understand what's on the AI horizon. And that includes the data centre industry. The whole supply chain is busy trying to assess the impact of AI on the data centre sector. Is it 'just another change', like virtualisation and cloud before it, or are we talking a seismic change which no one has quite yet understood?

Marc Marazzi sees the likely sheer pace and scale of AI adoption as a major challenge. "We are at the tip of the iceberg. There's the 'easy' customer service, the question and answer-based AI, which is already being adopted, but then, as applications become more sophisticated, we'll start seeing industries that we're not even thinking about yet adopting these types of high performance compute requirements. And that will drive a significant push on colocation to get applications up and running very quickly, because, as soon as an industry and its users start adopting AI, the proliferation of these tools will be unbelievable."

Marc cites the example of the Zoom application explosion when the pandemic hit. In December 2019, the firm had 10 million subscribers and by June 2020 this figure had risen to 200 million. The application and all the associated IT and data centre infrastructure had to scale up massively. And AI will dwarf such rapid expansion.

The challenge for the data centre is to be high performance ready. As Marc explains: "I guess it's going to be something of a balancing act because, sure, you can deliver it, but how efficient is that data centre going to be and how smart are you going to be in the investments you make so you don't need to change everything in 12 or 18 months' time? I think we'll see the owners and operators looking at technologies that are efficient, fast to deploy and easy to adapt if the requirements change. And Legrand technology really allows you to do that, it allows you to scale up without massive changes to your data centre infrastructure."

One of the major questions facing the industry is whether or not the AI workloads are best met through upgrading existing facilities, building brand new sites or perhaps a mix of the two (another hybrid approach).

Marc provides a real-world example of this dilemma: "I was recently in Oslo and I had such an interesting conversation with a CEO from a design and build company and a CEO from a colocation company. The common theme from them both was the big increase in power requirement and densities, which meant that they were having to look at data centres they'd built only 19 months ago and asking us to find ways of improving this infrastructure. People are looking at improved containment, improving airflow, or removing CRAC units and putting in rear door cooling. When it comes to new data centres, we are being asked for more efficient ways of deploying critical power. For example, our Starline track busway solution doesn't just remove the spaghetti junction of cables under the floor it also provides the ability to add tap-offs (plug-in units) very quickly and to change them as and when more functionality is required."

Marc continues: "We're also seeing a hybrid approach, where people will update and renovate parts of the data centre, to make way for AI or high performance computing when those clients come in. For Legrand, there's a major opportunity to help in terms of education – explaining what can be done. We're spending more time on the same side of the table, talking about how to specify and build a particular solution, as opposed to being given a shopping list from the customer, saying this is what we want."

Crucially, with the industry already on the road to sustainability, an AI build out cannot be allowed to derail this progress. So, the conversations Marc and his team are having with customers cover not just how quickly can a production site go into production so the end client can start delivering their services, but whether the technology solution is going to be sustainable. And this is a major strength for Legrand, which has always had very strong CSR goals. Whether through acquisitions, driving sustainability in the factories, or the way in which technology is deployed, the company thinks green as a matter of course.

> "We're also seeing a hybrid approach, where people will update and renovate parts of the data centre, to make way for AI or high performance computing when those clients come in. For Legrand, there's a major opportunity to help in terms of education – explaining what can be done"

Marc explains: "Our ColdLogik rear door coolers by USystems allow cooling up to 200 KW/rack using 800 watts of power for that particular rear door. Now, if we imagine what a CRAC unit costs to run to deliver that, it's incredibly inefficient in comparison to what a rear door cooler can achieve. And with the rear door coolers, you can start off small, so you're cooling, say, 72 racks, and if you need more you just add fans.

"If you built out a data centre with our Starline busways and busbars, our Minkels cabinets, which have the best airflow management on the market, and then the USystems rear door cooling, you have a very efficient data centre – and you will have freed up more space as well."

Legrand's AI/HPC data centre infrastructure portfolio

Marc has already explained how the ColdLogik RDHx rear door cooling solutions bring the cooling much closer to the cabinet, providing significant performance benefits when compared to CRAC units. Additionally, with these CRAC units taking

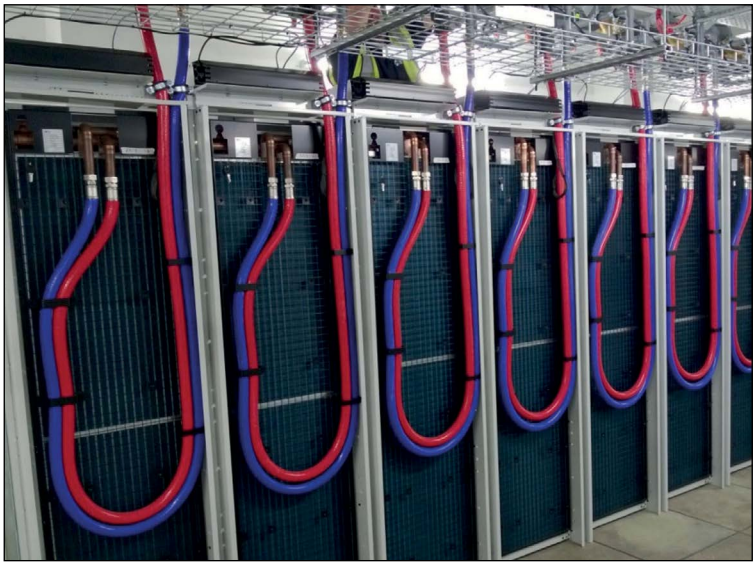

➤ USystems ColdLogik RDHX Rear Door Cooler

Raritan PX4 Intelligent Rack PDU

up to a third of the space in a data hall, the rear door technology allows this space to be reclaimed. For a colo, where each cabinet might be earning them £18,000 a year, the return on investment of installing the rear door cooling solution is extremely attractive.

The technology and commercial win-win is a similar story when it comes to Legrand's Starline busways, as Marc details: "With our busways providing continuous access you have the most effective way of implementing the power supply into the room – because it's all overhead. Imagine if you have Remote Power Panels (RPPs) either end and you've got the cabling going from these into each cabinet. If you have changes required in these cabinets, the data centre will typically not remove the cables that are already in there as they are worried about pulling out the wrong cable(s). So, they just add more cables. If you are making a number of moves, adds, and changes over the year, and you are not moving the cables, that's incredibly inefficient.

"You're not being kind to the planet as you are using a lot of cabling and you now have these huge RPPs. If you think there might be six to eight RPPs in a hall – think of the commercial benefit of removing them – and use the reclaimed space for racks."

One of the major benefits of the Starline busway solution is the custom nature of the tap-offs – you can have them in any configuration you require. Some clients buy certain types of tap-offs for specific applications that their customers might have, and also buy different tap-offs for other applications for the same customer. Engineers spend less time in the data centre because all they are doing is coming in and adding a tap-off. As Marc summarises: "Once you have the Starline busway in the data centre, you are just adding tap-offs for the amount of cabinets you need and the amount of changes that will be required."

Optimising power management is another key aspect of ensuring an efficient data centre. With AI and other HPC loads placing new, higher demands on the power supply, Legrand's latest intelligent rack PDU solutions are well placed to provide the necessary power monitoring and measuring. These PDUs provide real time visibility into the power consumption, enabling the operators to optimise energy use, prevent any overloads, and carry out predictive maintenance – looking at whether or not AI applications are going to cause any issues for any particular racks and/or servers. Intelligent load balancing is also a feature of the PDUs, which also offer energy saving modes.

The other PDU capability worth noting has to do with the outlets. Marc outlines: "Imagine someone rents out space in your cabinets and they are installing IT equipment with the old 'kettle plug' C13 connections. Their applications change – their project managers say, this is the IT equipment you should now use and these are C19 – but you are stuck with the old racks. Well, our intelligent rack PDUs have a hybrid outlet that can accommodate both the C13 or C19 connections."

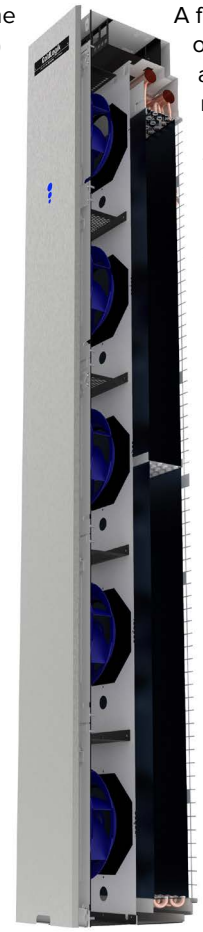

A familiar theme with the PDUs, as with Legrand's other data centre infrastructure, is its flexibility to adapt when equipment moves and changes are made. Additionally, they can be colour coded, so users can work out the different phases they have got for that type of load balancing. Maximum efficiency and flexibility for the user.

Another solution of Legrand's data centre infrastructure portfolio is the Minkels Nexpand cabinet platform. Marc takes up the story: "Minkels really spent a lot of time on airflow management – especially on the cable entry and exit points. And we continue this work, which has come full circle because airflow management, being able to secure and understand exactly where your air is going, has again become super important from an efficiency standpoint. Also important is the custom nature of the platform. We don't tell customers 'these are the racks that we've got, which one do you want?', but we ask them what they need.

"We're definitely seeing the need for more high density racks and Legrand is putting in extensive engineering effort to develop the Nexpand platform, focusing on reducing energy consumption and improving the overall performance. They work very closely with our PDU factory, so the way that the PDUs are mounted, and providing enough space for

people to work within the rack, all these things are designed to help customers get their production sites up and running as quickly and efficiently as possible."

Legrand provides an integrated data centre focus

Marc summarises the breadth and depth of Legrand's data centre infrastructure portfolio: "We have the high efficiency transformers, the switchgear, the UPS and we have the high performance busbars that connect up these solutions. And from the UPS, we have the Starline busway going into the data centre, the cable management going into the data centre and dropping down into the Minkels cabinets, with the intelligent PDUs, and we even provide remote access to the critical devices as well. That's a full suite of products, all from Legrand."

Marc continues: "And this was the genesis of the idea behind the creation of Legrand Data Center Solutions in 2021. We've got all these solutions, let's get the teams, the engineers, the factories all working together so we can provide a total solution and the consultancy to go with it."

Given the challenges facing the data centre sector at the current time, and with the AI 'effect' yet to be felt, this last aspect – consultancy – is playing an increasingly important role for Legrand as it works alongside its customer base. Marc explains: "We're working with both consultants and end users when it comes to legislation, for example, helping them determine certain elements of the data centre infrastructure and how they should approach it. We'll provide comprehensive assessments of existing infrastructure to help identify potential bottlenecks, areas which they should be thinking about. And, of course, we advise on what we're seeing with AI workloads and how to build out the data centre to meet this demand."

Legrand also provides Innovation Days. Attendees may well know what a PDU is, and what a cabinet is, but, via the Innovation Days, Legrand provides a different way of looking at the components within the data centre and explaining the innovation possibilities within its solutions that customers maybe haven't fully thought about. Continuous Professional Development (CPD) days are also being rolled out this year, helping attendees to understand best practices regarding data centre infrastructure, design, installation, and operation.

While Marc is still acclimatising to his new role heading up the Legrand Data Centre Solutions business (still learning before the leading phase, as he puts it!), he has identified one area where he can see a major benefit to customers. "If there's one way in which I think we can really help our customers it's the turnaround speed in getting data centre designs to them, and getting this information to them in an easier to understand format," Marc says. "Internally, we have these dashboards that tell us where orders are and, when they are going to ship. But I would like us to have more of an interactive way that clients can obtain information from us much quicker. "We almost need our own AI which will provide our customers with the answers they need! The requirements are going up so quickly that I want to make sure we're servicing the customers and providing operational excellence. And this comes down to streamlining our processes, being very efficient, providing the customers with the best experience."

➤ Starline Track Busway

Learn more about how Legrand can prepare your business for AI by visiting its website here.

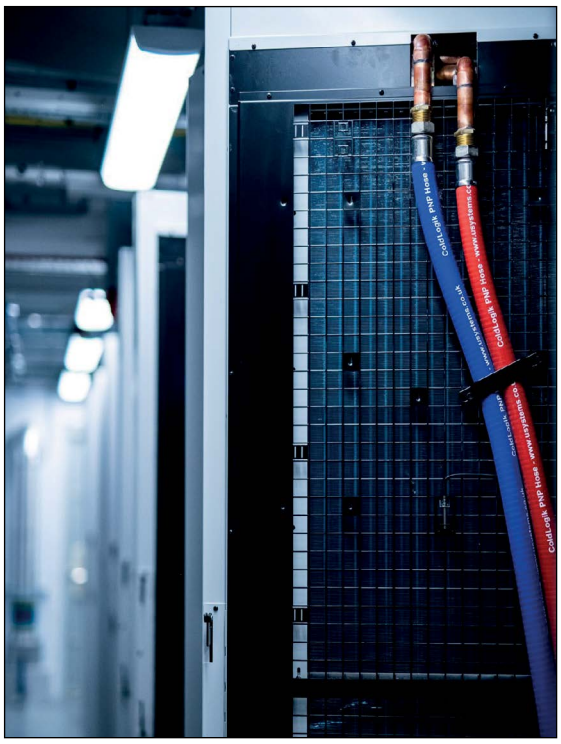

➤ USystems ColdLogik RDHX Rear Door Cooler

Is the data centre industry ready for AI?

STEVEN CARLINI, VP FOR INNOVATION AND DATA CENTRE AT SCHNEIDER ELECTRIC, explains how the massive, continuing AI explosion is threatening to re-write the rules when it comes to data centre design and operation. The company recently published a Blueprint for AI data centres, which explores pioneering approaches to data centre design, and Steve shares the four key AI attributes and trends that will underpin physical infrastructure requirements in the era of machine learning.

THE MASSIVE, ongoing AI explosion promises to have a major impact on the data centre. But does this impact require some slight adjustments, so business pretty much as normal, only much more of it, or will it potentially mean the rulebook needs to be thrown out and data centre design will need to be 'reinvented'?

Steve responds: "There is an ongoing, insatiable capacity for data centres. But these new data centres, especially the AI training data centres, are much, much different than the data centres that we have today - the X86 type of server farms that are massive in scale, but which tend to be lower density and spread out more, have a larger footprint. The new servers with the GPU accelerators, or ASICs, are much different. They require a lot more power. All of the processors run in parallel, so you package these in servers, but they're not acting as an individual server, they're acting as one giant server. So, you could have 10,000 GPU accelerators operating in parallel in these data centres. So it is, from a training perspective, a need for an incredible amount of capacity because of all the parameters that we're putting into these to train the models and also the different formats of the parameters."

He adds: "At the beginning it was just text, but now you're feeding them with images and even videos, and they're outputting different formats as well. So multimodal AI is here and it's much different from a data centre perspective, from a capacity standpoint - very high capacity, the need for a lot of data centres, so very high scale and also very high density."

The obvious place to start to meet this AI demand would be to upgrade existing legacy data centres, but is this economically viable and/or practical? Steve says that it is possible to retrofit existing sites, but if you don't have a lot of reserve power or excess power feeding that data centre, they're going to, number one, have to source more power, and number two, have to densify that power. In other words, running all of the power to a smaller, smaller area. So, when you're thinking of the power train and you're thinking of bringing it down from medium voltage, which can be 10,000 to 39,000 volts, down to the server, it's optimal to reduce the number of transformations that you have. Therefore, it is more economical to start from scratch and then have your powertrain optimised for that layout.

We're not talking about cables anymore, we're talking about a very high power busway, and then distributing that to the servers is also very complicated, so there's a tremendous amount of work involved, and that's just the power. We haven't even started to talk about the cooling. So, a legacy data centre could be adapted for AI workloads, but you could potentially end up with a lot of excess floor space in these data centres. Each hall would have one little corner of AI, high density computing, or cluster AI.

Bearing in mind the problems converting legacy data centres to be AI-capable, and the number of data centre owners and operators who see the opportunity and want to be able to tell customers that they can accept AI workloads, what can end users expect from their data centre provider? In other words, if a provider is saying that their facility

is perfectly fitted for a customer's AI application, what might they actually mean?

Steve comments: "From a physical infrastructure perspective, you see the companies that have been doing it for a little while, a lot of them actually paused construction until they could figure out exactly how to build them at the correct density. But I think you're talking about maybe some of the service providers that are retrofitting some of their sites and building cages or clusters to accommodate these AI servers. And in most cases, and we've been dealing with a lot of these companies, they are designing their cages to be able to bring the appropriate amount of power and the appropriate amount of cooling to those cages. Or if they're building the cluster themselves with the racks, they could build out the entire system with the servers, or build out the power distribution to the racks and the cooling and the piping and the manifolds all the way to the racks as well. And you could put in your own servers. But we're seeing a big move.

There's, like I said, insatiable demand for this capacity. So, they are retrofitting many of the sites and we've been working with them for years on this and it never really was very slow going! Now, all of a sudden it's been like a light switch going on in terms of the demand.

"The other thing I should mention which is interesting is the traditional data centre workloads were variable. So, based on your business and based on the amount of processing and storage that needed to be done, the power levels would go up and down. When you start AI training - and I'm still talking about training models - when you're training a model, a large language model, your servers are on in this parallel design and they're running at capacity or very close to capacity. You have to be very careful to design these. If you design it for 100 racks and you put in 100 racks of IT, it's going to run at 100 the entire time. There's no built-in buffer like the old days!"

Talking of which, with data centres already under pressure to source ever more power just to meet the demands of existing digital transformation, let alone the 'new' AI boom. And, of course, the power needs to be renewable.

There has been a lot of work within the industry to make the data centres as efficient as possible and run the data centres with the most renewable sources as possible. And for every data centre that's being permitted or being constructed after the permit, there has to be an identified source of renewable power for these sites. And an increasing number of countries are getting very stingy with their permitting. It's a huge challenge and it started a couple of years ago and last year was not as bad as a lot of people predicted, with the plan brownouts and planned blackouts. Nevertheless, there is a very real concern.

There are a couple of issues. The renewable capacity that was planned to be built, there were different reasons why it didn't get built. There were some legal issues, there were some capacity issues with offshore wind, with the number of rigs that they had to be able to put them in, and then the whole supply chain. Wind and solar has not favourable for the last couple of years, but there is now a big acceleration in the amount of renewables this year. And there are a lot of government programmes and the data centre industry has been funding a lot of these projects directly or indirectly through PPAs and energy credits and the like.

From a technology perspective, there's going to have to be a bigger focus on energy storage, either on the grid or at these data centre sites. If there's a situation, for example, where you're running primarily on renewables and there's excess renewables, like during the day when the sun is shining, the wind is blowing, you would want to take advantage, instead of curtailing that power, to put energy storage on your site and put all that power into batteries where you could better utilise it or give it back to the utility when they need it. As a result, the whole grid is becoming much more distributed and much more complicated, and there's going to be new and innovative ways to manage the grid.

Keeping cool

With the increased power demand for the increased compute loads comes whole new levels of heat being generated. Which begs the question as to whether or not liquid cooling will have an increased role to play. Does the whole industry have to move in this direction or does air cooling still have a significant role to play? And what's the best AI cooling?

Steve explains: "If you look at the biggest models that have been trained, that are in production today, they were all trained primarily on air cooled servers. So, it is possible to use air cooled servers, but from a sustainability and efficiency perspective, it's not optimal. And from a cost perspective, it's not optimal

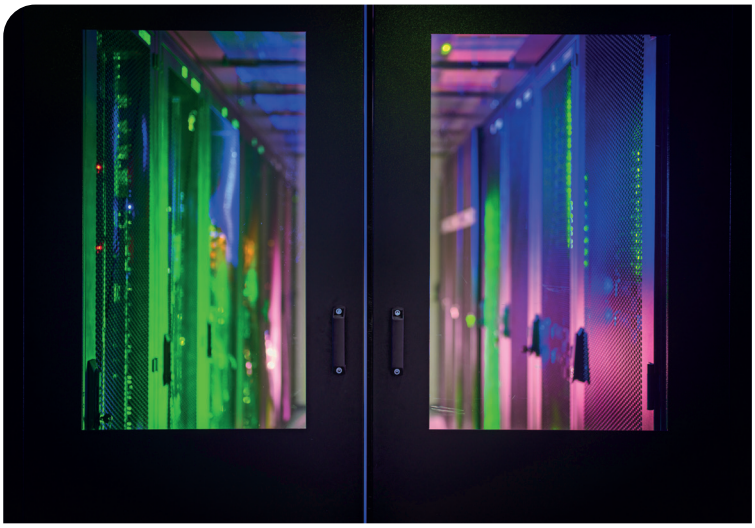

because of the more you spread out the load, the more expensive it gets, the more real estate that you take up. The current designs, the designs that are coming up, are going to be more densified, so the advantages are going to be they are going to be more sustainable and lower cost, so you can do it with air. But it's a funny looking data centre where you have a big IT rack and one or two servers in each one and it's all spread out. And in terms of the Infiniband and the fibre, you can't run 20 gig or 40 gig connectivity, you're going to be running 400 or 800, which is much more expensive because like I said, these are all running in parallel, so the connectivity between all of them has got to be lightning fast."

When you deploy in the field, the AI models don't have to be connected back to the core data centre, because once the model is trained, you're going to deploy a subset or a compressed version of that model in the field as an inference model

One of the potential pluses of moving to the liquid cooling option is the opportunity for waste heat reuse. Which begs the questions, if someone is building a new data centre with AI loads in mind, will they try and build it in a location where the heat transfer is easy to achieve? And/or, will government at whatever level demand such waste heat reuse? Waste heat reuse is seen a lot more in Europe than in the US and in parts of Asia, whether through ither voluntary or mandatory standards, like Germany just initiated a standard that 20% of the weight of the heat had to be reused. Everyone talks about district heating, which is heating of office buildings or homes, and because of the AI training models running, 100% of the time, the heat is more constant. In one example, Schneider Electric built a data centre where it was next to a plant that did pellets, the pellets for pellet stoves. These were made from either wood chips or excess newspapers, back when we had newspapers, but they would dry out these things and compress them into pellets, using data centre waste heat. Those types of applications are ideal because you don't really need a consistent heat all the time. At the same level, it could vary a little bit and not have a dramatic effect. Overall, there is more and more planning of how to position the data centres next to areas where the heat could be reused.

Data centre location is also important in terms of taking advantage of favourable climatic conditions (ie naturally cool/cold temperatures) and a plentiful supply of renewable energy sources. Might we end up with AI factories in these favourable geographical locations – the Nordics for example – because it makes more sense financially and sustainably, or is latency still an issue?

Steve outlines: "Everyone's familiar with kind of the cloud computing model, where you have kind of the core centralised data centre, regional data centres, and then edge data centres, and they're all sharing data, they're all sharing control planes. There are latency issues when these models are just churning and they're not really connected and the data doesn't need to be duplicated everywhere. I like to talk about integrating the supply chain and dynamically building based on the information from other IT systems, but when you're using that, it has to be more real time, so you can't have this big lag like when you're writing a book report. You could have those data centres located somewhere else and you could wait five minutes. But if you're running grid management too, or grid automation to supply power, those need to be real time as well. "When you deploy in the field, the AI models don't have to be connected back to the core data centre, because once the model is trained, you're going to deploy a subset or a compressed version of that model in the field as an inference model. And based on how fast you need the response or the decision, and based on how accurate you need the decision and how comprehensive it needs to be, that would dictate the size of the model that you need to deploy. So, you have these giant LLM models and you're going to deploy, in some cases, a model that's only 5% of the size of that in the field, and it doesn't need to be connected to the core data centre at that point either."

Higher densities, high data centre impact

We've covered power and we've covered cooling related to the new AI workloads But there is also the not so small matter of the cabinets and racks. If we're looking at higher densities, is there a need to re-engineer data centre infrastructure solutions to make them more robust?

The HPS servers are much larger, much heavier and deeper. As a result normal racks have already been stressed, and in a lot of cases, there's going to be custom racks that are higher, that can hold a lot more weight. A recent Schneider Electric White Paper includes some of the recommendations for how much weight it's going to need to hold, how deep it's going to need to be, and it's going to put tremendous pressure on existing infrastructure. For example, you either like raised floors or you don't. Thew dense workloads make it really challenging to deploy in raised floor environments. On the software side, there's an incredible amount of software that's looking at the GPUs and the health of the GPUs and the speed, so there's a tremendous amount of focus on the computational side of these AI models. So there needs to be an equal amount of focus on the physical infrastructure supporting these models as they're training and as they're running, especially as organisations start deploying more and more in the field and at the edge. Edge computing has been talked about for years and it may well be that AI is actually going to be one of the catalysts

where we're going to see much faster and larger scale deployment at the edge.

How is the industry responding to the AI challenge?

Steve is confident that the data centre industry is ready to meet the AI demand. He says: " Right now the industry is in a very steep uptick and is responding by adding a lot of capacity and adding the capacity to support these types of loads and some of the service providers are adapting to change and to support these type of loads as well. The stress at the moment is really for the enterprise that wants to run some of their legacy applications in cloud environments. A lot of the colo environments and a lot of the colos are transitioning to support the AI loads and it's really causing an economic supply and demand issue where you may have had your application in your colo for the past five years, and now all of a sudden there's such demand for that space that your next contract negotiation may be a lot different than it was in the past."

AI also present a challenge in terms of sustainability. Many, if not all, of the companies that are deploying AI workloads do have carbon neutral commitments and a lot of them are on the road to net zero. There is a tremendous scale out. Data centres are becoming are more power hungry, but they're doing it in a way where they're offsetting the increase in capacity with different vehicles, whether they're credits or PPAs or building the renewable supply themselves. Data centre operators are investing in a lot of wind farms and solar farms, and also they're offsetting water usage as well. They have to meet their carbon neutral and water neutral supply targets, to be waste neutral. As a result, there is a tremendous amount of focus on building these data centres in a sustainable way. Schneider Electric is projecting that, even though the capacity over the next ten years is going to double, the carbon emissions are going to be relatively flat. Data centres are going to scale and scale very quickly, but the amount of carbon emissions is not going to scale with it.

It is also worth bearing in mind that data centres are being used by virtually all industry sectors to combat their own carbon emissions as well. For example, coordinating transportation, logistics and charging with the most renewable sources and also even the carbon output. A lot of these companies are using data centres and AI to actually reduce their carbon footprint. So, even though there's more data centre capacity being deployed, some of that's being used to fight carbon emissions in other industries as well.

Schneider's role

Schneider Electric itself is focused on data centres of all capacity. It is helping with the largest providers of AI data centre capacity, the internet giants. Helping them with their power systems, helping them with the software management of those power systems and of their cooling control. The company is also involved with the service providers, the

colo providers and the new machine learning as a service providers that are emerging. There's a new category of data centre operators that are just building AI training data centres with the intention of having other companies come to them to train their models. Schneider Electric is putting the solutions together to be able to enable the new types of AI servers to be deployed from a power perspective, from a cooling perspective and from a management perspective. Liquid cooling is definitely desirable from a sustainability perspective, from an efficiency perspective - even though it's been around for a long time, it's quite immature with the number of deployments to date.

In summary, Schneider Electric is helping a lot of companies with the whole idea of reimagining what the data centre is going to look like for these environments is something that we're working on. It is working across the board as a partner and as a solution provider as well, to be able to deliver the power equipment, the cooling commitment, the equipment for these types of applications.

Will AI break the industry?!

A question that has to be asked. If everyone does AI as much as it's perhaps predicted, at some point the sign will go up – data centres are full, no more capacity. Or will the industry always come up with the right answers?

Steve replies: "I think it has the potential to break the industry, but there are certain controls, like the amount of available power, the amount of available renewable power are very limiting. The fact that a lot of countries are requiring a lot more diligence for permitting is another. And then the whole supply chain issue. The equipment providers like Schneider, we can build so much equipment, so it is almost natural kind of regulation to how fast these things are going to be deployed. I was at conference a few weeks ago and there were a couple of speakers that said if we could deliver 20,000 of these data centres tomorrow, we would have customers for them. There is this insatiable demand and the industry is reacting to deliver it as fast as it possibly can. But I don't think it will break the industry. I think it'll be more of a controlled rollout based on the factors that I have talked about."

Designing sustainable data centres is key to preparing for a more digitalised and greener future

Building a greener future has become a key priority for individuals and businesses in the UK. Yet, 78% of the energy we use comes from fossil fuels which are responsible for most of the UK's territorial emissions of climate-changing greenhouse gases. Using energy more efficiently and intelligently in everyday life is vital to averting climate change.

By Dave Archer, National M&E Sales Manager at Mitsubishi Electric

BUT IN TODAY'S society, where technology is evolving rapidly and everything is digitalised, minimising energy consumption can be a real struggle as more digital information is being created and stored in huge, energy-intensive data centres. Making data centres truly sustainable by controlling energy use and preventing waste must be a priority to prepare for the future, though it isn't a straightforward issue.

The challenges raised by data centres

Data centres are becoming increasingly important as the amount of data generated in all walks of life increases. This means the sector's energy efficiency is under increasing scrutiny. Data centres are currently responsible for 1.5% of global energy consumption, and we expect it to reach 8% in 2030. As data centre tech becomes ever-more supercharged, finding ways to conserve energy is vital. But although the industry has worked hard to increase computing capacity while controlling energy use, improving energy efficiency is not straightforward. Sustainability is not always taken into account when designing data centres.

We can't prevent data centres from consuming more energy to produce more power in the future. Data shows the largest data centres produce over

100MW of energy, and we anticipate the increasing adoption of Generative AI to double or triple energy consumption in data centres. However, addressing the rising concerns about sustainability must be a core area of focus for designers and data centre operators. With change coming rapidly, they must consider all the options available to improve energy efficiency in data centres – this will be key to achieving the industry's goal of reaching net-zero by 2030.

Selecting the right cooling solution

It is key for businesses to ensure their IT systems allow them to quickly and easily access the data they need. The shortest outages can be extremely disruptive and greatly impact productivity and revenue. Therefore, IT servers must be kept in optimal conditions, which means adopting a reliable cooling solution.

Finding areas for conserving energy is vital, and data centre cooling systems are a potential source of savings, with up to 40% of total energy consumption in a data centre coming from cooling systems. Adopting the right cooling solutions in data centres not only boosts productivity by reducing the risk of an outage but also reduces the carbon footprint, enables the use of renewable energy sources and minimises operational costs. Data centre operators aim to upgrade facility cooling systems to transition to low carbon, cost-effective solutions and future-proof against rising heat requirements.

For example, Mitsubishi Electric offers DX Computer Room Air Conditioning (CRAC) solutions that control temperature and humidity in small and medium-sized enterprise data centres. They create an appropriate environment for IT systems while reducing energy consumption and running costs.

Re-using heat to avoid wasting energy

With great power comes more heat. Data centres eject a lot of heat, and one of the most critical questions for the future will be what to do with it. What is certain is that this heat shouldn't go to waste. This is an environmental and financial question – wasted heat means wasted energy, unnecessary costs and a significant impact on the planet.

One development gaining popularity is re-using the heat from data centre servers to heat homes and buildings and cut CO2 emissions. We expect the data-centre heating market to be worth $2.5 billion by 2025. The heat removed from the data centre servers can provide hot water to other buildings. Some high-temperature systems can also directly heat other buildings where appropriate. The temperature doesn't have to be high – low-temperature rejected heat can be combined with heat pumps to raise the water temperature for domestic hot water and space heating in offices and homes.

If we treat the ejected heat as a form of energy, then the combination of efficient cooling and heat pumps for modern ambient networks can make all the difference. This approach, known as an ambient heat loop or Fifth Generation heat network, saves energy and costs. For example, Dutch company Switch Datacenters has replaced its gas generator units with data-centre heating to reduce its reliance on natural gas. Eventually, the organisation delivered 97% of its server heat to homes and offices to improve energy efficiency and allow customers to save power costs.

There is little doubt that the modern world needs data centres. Today's increasing flow of digital data needs to be stored securely and in optimal conditions, but not at the price of the environment. Cutting gas emissions and energy consumption must remain among the top priorities. Harnessing energy-efficient cooling by re-using ejected heat is a practical solution to minimise waste, reduce emissions and contribute to building a more sustainable and modern future.

There is little doubt that the modern world needs data centres. Today's increasing flow of digital data needs to be stored securely and in optimal conditions, but not at the price of the environment

Increasing data centre power consumption in an energy crisis

The role of generator maintenance in ensuring power resilience.

BY CRESTCHIC

THE GLOBAL ENERGY CRISIS is affecting the profitability and operational efficiency of businesses from all sectors, but for data centres and their inherently higher power consumption levels, the impacts are polarised. For the data centre market to achieve its projected growth trajectory, it must continue to adopt power-intensive next-generation IT technologies. And, it must do that in the face of an unstable grid, potential planned blackouts and suggested power rationing. That means rigorous investment in backup power and generator maintenance, at a time when energy bills are consuming investable profit. Paul Brickman, Commercial Director, for Crestchic Loadbanks, explores the often-overlooked necessity of loadbank testing in securing data centre power in an energy crisis.

Hiked power prices as a result of Russia's fuel sanctions, and the threat of planned blackouts or

power rationing to better manage an unstable grid, are forcing data centres to redistribute investment away from key business functions like training and recruitment, and towards improving energy efficiency and securing power resilience.

The current geopolitical challenges and their far-reaching implications are recognised in the 2023 Uptime Institute Report – Five Data Center Predictions for 2023 – in which the authors acknowledge that the Russian fuel sanctions, along with technological challenges "will make the planning of data centre development and operation more difficult."

Powering next-generation data centres

According to the Uptime Institute, IT hardware has been fairly standard in terms of its draw on mainstream server capacity for a few decades now, creating technical stability and relatively constant power and cooling requirement. This has enabled data centre designers to accommodate several IT refreshes without major upgrades to server technology and the associated hike in power consumption that would come with it. Essentially offering the latest technologies, without an increase in energy requirement.

This grace period is now over. Power requirements for next-generation IT hardware are far higher, rack power density is increasing, and "hotter" processors are putting pressure on the performance parameters of existing infrastructure. This rapid rise in IT power density means server power consumption is on a steep climb, creating a need for more power, at a time when the grid is unstable, and energy prices are at their highest in decades.

Extreme-density racks are now commonplace in technical computing too, as well as high-performance analytics and artificial intelligence training. If data centre operators want to successfully penetrate these niche markets, they will need to foot the heightened energy bill and maintain a robust power resilience regime until the situation stabilises.

Data centres take action

Demand for next-generation IT technologies will not slow. As a result, data centres are forging ahead with the essential upgrades required to UPS, batteries, switchgear and generators to accommodate increased power density.

Extreme-density racks are now commonplace in technical computing too, as well as high-performance analytics and artificial intelligence training

This is a risk. With an unstable grid and planned blackouts still on the agenda for many governments, upgrades will need to be bolstered with a watertight energy resilience plan to protect against power fluctuations and total outages.

The importance of a power resilience strategy in an energy crisis

According to research from the Ponemon Institute, in its third "Cost of Data Center Outages" report, the total cost of downtime has continued to rise over the last six years – rising a staggering 38% to $740,357 per incident. That equates to nearly $9,000 per minute – a figure that many data centres will have to swallow if the grid becomes too unstable to sustain their increasing power consumption.

Many data centres will already have a robust generator testing and maintenance regime in place, but the use of load banks is often overlooked, especially when budgets are stretched.

At a time when power outages are more likely, load bank testing should play an integral role in a data centre's energy resilience strategy and it would be prudent to evaluate your strategy in line with the current landscape to ensure it is robust enough.

With this in mind, what would be considered best practice for testing a backup power system?

Best practice load bank testing in times of crisis

Ideally, as a minimum all generators should be tested annually for real-world emergency conditions using a resistive-reactive 0.8pf load bank. Best practice dictates that all gensets (where there are multiple) should be run in a synchronised state, ideally for 8 hours but for a minimum of 3.

Where a resistive-only load bank is used, testing should be increased to 2-4 times per year at 3 hours per test. In carrying out this testing and maintenance, fuel, exhaust and cooling systems are effectively tested and system issues can be uncovered in a safe, controlled manner without the cost of major failure or unplanned downtime.

The alternator is not thoroughly tested though, with a resistive-only test and therefore a resistive-reactive test would always be recommended. It may be advisable to test more frequently during times of crisis for added peace of mind.

Why is resistive-reactive the best approach?

Capable of testing both resistive and reactive loads, this type of load bank provides a much clearer picture of how well an entire system will withstand changes in load patterns while experiencing the level of power that would typically be encountered under real operational conditions.

Furthermore, the inductive loads used in resistive/reactive testing will show how a system will cope with a voltage drop in its regulator. This is particularly important in any application which requires generators to be operated in parallel (prevalent in larger business infrastructures such as hyperscale data centres) where a problem with one generator could prevent other system generators from working properly or even failing to operate entirely.

This is something which is simply not achievable with resistive-only testing.

Navigating growth when power is scarce

No matter the geopolitical challenges and the effect it is having on power availability, data centres have no choice but to grow.

Demand will not cease, and power-intensive next-generation technologies are unavoidable. Ensuring power resilience via a watertight backup power supply and a robust testing and maintenance regime will enable data centre designers and operators to grow, safe in the knowledge that, should planned blackouts, power rationing or grid fluctuations happen, the power will always remain on.

For support with improving your load bank testing regime, search Crestchic Load Banks online.

AFL at DCW: Visit us at stand D825

Join us to discuss your hyperscale data center plans for 2024 and beyond. Learn more about our innovative fiber connectivity solutions and explore how we can help you adapt and scale to meet the evolving needs of your growing network.

> Gain valuable insights.
> Unlock next-level Modular Connectivity.
> Overcome hyperscale data center challenges.

Scan to get your tickets now

Why hybrid cooling is the future for data centre

Rising rack and power densities are driving significant interest in liquid cooling for many reasons. Yet, the suggestion that one size fits all ignores one of the most fundamental aspects of potentially hindering adoption - that many data centre applications will continue to utilize air as the most efficient and cost-effective solution for their cooling requirements. The future is undoubtedly hybrid, and by using air cooling, containment, and liquid cooling together, owners and operators can optimise and future-proof their data centre environments.

BY GORDON JOHNSON, SENIOR CFD MANAGER, SUBZERO ENGINEERING

TODAY, many data centres are experiencing increasing power density per IT rack, rising to levels that just a few years ago seemed extreme and out of reach, but today are considered both common and typical while simultaneously deploying air cooling. In 2020 for example, the Uptime Institute found that due to compute-intensive workloads, racks with densities of 20 kW and higher are becoming a reality for many data centres.

This increase has left data centre stakeholders wondering if air-cooled IT equipment (ITE), along with containment used to separate the cold supply air from the hot exhaust air, has finally reached its limits and if liquid cooling is the long-term solution. The answer is not as simple as yes or no, however.

Moving forward it's expected that data centres will transition from 100% air cooling to a hybrid model encompassing air and liquid-cooled solutions with all new and existing air-cooled data centres requiring containment to improve efficiency, performance, and sustainability. Additionally, those moving to liquid cooling may still require containment to support their mission-critical applications, depending on the type of server technology deployed.

One might ask why the debate of air versus liquid cooling is such a hot topic in the industry right now? To answer this question, we need to understand what's driving the need for liquid cooling, the other options, and how can we evaluate these options while continuing to utilize air as the primary cooling mechanism.

Can Air and Liquid Cooling Coexist?

For those who are newer to the industry, this is a position we've been in before, with air and liquid cooling successfully coexisting while removing substantial amounts of heat via intra-board air-to-water heat exchangers. This process continued until the industry shifted primarily to CMOS technology in the 1990s, and we've been using air cooling in our data centres ever since.

With air being the primary source used to cool data centres, ASHRAE (American Society of Heating, Refrigeration, and Air Conditioning Engineers) has worked towards making this technology as efficient and sustainable as possible.

Since 2004, its published a common set of criteria for cooling IT servers with the participation of ITE and cooling system manufacturers entitled "TC9.9 Thermal Guidelines for Data Processing Environments".

ASHRAE has focused on the efficiency and reliability of cooling the ITE in the data centre. Several revisions have been published with the latest being released in 2021 (revision 5). This latest generation TC9.9 highlights a new class of high-density air-cooled ITE (H1 class) which focuses more on cooling high-density servers and racks with a trade-off in terms of energy efficiency due to lower cooling supply air temperatures recommended to cool the ITE.

As to the question of whether or not air and liquid cooling can coexist in the data centre white space, it's done so for decades already, and moving forward, many experts expect to see these two cooling technologies coexisting for years to come.

What Do Server Power Trends Reveal?

It's easy to assume that when it comes to cooling, a one-size will fit all in terms of power and cooling consumption, both now and in the future, but that's not accurate. It's more important to focus on the actual workload for the data centre that we're designing or operating.

In the past, a common assumption with air cooling was that once you went above 25 kW per rack it was time to transition to liquid cooling. But the industry has made some changes in regards to this, enabling data centres to cool up to and even exceed 35 kW per rack with traditional air cooling.

Scientific data centres, which include largely GPU-driven applications like machine learning AI and high analytics like crypto mining, are the areas of the industry that typically are transitioning or moving towards liquid cooling. But if you look at some other workloads like the cloud and most businesses, the growth rate is rising but it still makes sense for air cooling in terms of cost. The key is to look at this issue from a business perspective, what are we trying to accomplish with each data centre?

What's Driving Server Power Growth?

Up to around 2010 businesses utilized single-core processors, but once available, they transitioned to multi-core processors, however, there still was a relatively flat power consumption with these dual and quad-core processors. This enabled server manufacturers to concentrate on lower airflow rates for cooling ITE, which resulted in better overall efficiency.

Around 2018, with the size of these processors continually shrinking, higher multi-core processors became the norm and with these reaching their performance limits, the only way to continue to achieve the new levels of performance by compute-intensive applications is by increasing power consumption. Server manufacturers have been packing in as much as they can to servers, but because of CPU power consumption, in some cases, data centres were having difficulty removing the heat with air cooling., creating a need for alternative cooling solutions, such as liquid.

Server manufacturers have also been increasing the temperature delta across servers for several years now, which again has been great for efficiency since the higher the temperature delta the less airflow that's needed to remove the heat. However, server manufacturers are, in turn, reaching their limits, resulting in data centre operators having to increase the airflow to cool high-density servers and to keep up with increasing power consumption.

Additional Options For Air Cooling

Thankfully, there are several approaches the industry is embracing to cool power densities up to and even greater than 35 kW per rack successfully, often with traditional air cooling. These options start with deploying either cold or hot aisle containment. If no containment is used typically, rack densities should be no higher than 5 kW per rack, with additional supply airflow needed to compensate for recirculation air and hot spots.

What about lowering temperatures? In 2021, ASHRAE released their 5th generation TC9.9 which highlighted a new class of High-Density Air-Cooled IT equipment, which will need to use more restrictive supply temperatures than the previous class of servers.

At some point, high-density servers and racks will also need to transition from air to liquid cooling, especially with CPUs and GPUs expected to exceed 500 watts per processor or higher in the next few years. But this transition is not automatic and isn't going to be for everyone.

Liquid cooling is not going to be the ideal solution or remedy for all future cooling requirements. Instead, the selection of liquid cooling instead of air cooling has to do with a variety of factors, including specific location, climate (temperature/humidity), power densities, workloads, efficiency, performance, heat reuse, and physical space available.

This highlights the need for data centre stakeholders to take a holistic approach to cooling their critical systems. It will not and should not be an approach where we're considering only air or only liquid cooling moving forward. Instead, the key is to understand the trade-offs of each cooling technology and deploy only what makes the most sense for the application.

How augmented reality is impacting operational performance and total cost of ownership

Augmented Reality (AR) is gathering momentum in the world of equipment servicing and is fast becoming the new norm, impacting the way in which shopfloor operatives, service/maintenance technicians and engineers interact. The rise in popularity for AR and how it is helping maintain competitiveness in an ever changing and demanding economic climate.

By Anna Mazzoleni, Service 4.0 Global Product Manager for ABB Electrification Service

Augmented Reality (AR) is gathering momentum in the world of equipment servicing and is fast becoming the new norm, impacting the way in which shopfloor operatives, service technicians and engineers interact.

Here Anna Mazzoleni, Service 4.0 Global Product Manager for ABB Electrification Service, discusses the rise in popularity for AR and how it is helping maintain competitiveness in an ever changing and demanding economic climate.

Research shows that businesses are concerned about the impacts of energy pricing and security and the impact that this is having on their ability to compete, invest in people and reach their sustainability targets.

This is according to ABB's Energy Insights Survey of 2,300 leaders from small and large businesses, which revealed that 92 percent of respondents believe the continuing instability of energy is threatening their profitability and competitiveness. It is also having a significant impact on the workforce with decreased investment in employees. They are also concerned about potential impacts of meeting their sustainability targets.

The priority for industry, therefore, is resilience! Operators need to build an infrastructure that is resilient to all geopolitical changes and be willing to adopt and integrate new technologies, such as AR.

For example, the likes of augmented servicing guides visible on the electrical equipment, boosts interaction and the ability for self-support, removing the need for an engineer to travel to the facility, which increases efficiency, reduces downtime and drastically cuts CO2 emissions.

So, what is AR?

Unlike Virtual Reality (VR), which replaces physical reality with a computer-generated environment, AR superimposes digital information on the physical world.

Through AR technology, operational information is presented in a completely new way – augmented in a person's view of their real environment and acting as a digital assistant. But crucially, AR makes digital

assistance interactive, more practical to absorb, as well as easier to understand and act upon. In other words, its core capabilities are visualize, instruct and interact.

The technology has been around since the early 1990s, but it wasn't until the outbreak of the COVID-19 pandemic that it was more widely adopted. With international and domestic travel bans in place, operators needed to find alternative ways to carry out essential servicing and maintenance. And that's where the quantifiable benefits of digital support technology using AR proved compelling. Indeed, AR has unlocked post-pandemic productivity by simplifying maintenance, reducing downtime costs and increasing equipment effectiveness via the quality

of repairs and speed of resolution. It has also reduced the need to travel to site, but with an expert always to hand, albeit remotely. Through innovation, we are now connecting remote live experts to 'real life' customer issues, wherever they are in the world, to reshape the way we interact with them and enable self-learning through first-rate support.

AR also delivers value; making procedures faster, smarter and safer, in a standardized way to facilitate knowledge retention and continuous improvement cycles. For example, our remote factory acceptance testing (FAT) now regularly deploys augmented reality solutions encompassing audio, video, document sharing and live annotations by overlaying digital information onto the equipment operatives are working on, removing the need for customers to visit our facility.

Solutions in hours not days

Because AR applications work on a multitude of devices, it is no longer a restricted technology either. Even the Android and iOS mobile devices we use in our daily lives can provide the operational gateway to reducing downtime and increased efficiency. The service expert gets real-time visual insight to the application, accessing chat, images or videos shared by the on-site engineer, and in turn ABB helps the customer troubleshoot by guiding through the service process with the aid of interactive tools that visualize and simplify the instructions.

Seeing is believing, of course and some recent outcomes show the value of AR in practice. Together with ABB experts the field service engineers of a leading pulp and paper producer now use Microsoft Hololens headsets imparted with AR technology containing repair strategies and guidance documentation – so that maintenance issues that would normally result in days of downtime for travel, troubleshooting and resolution are instead solved in hours.

Additionally, one of the world's largest marine shipping operators needed remote maintenance to support problem-solving for its global fleet and reduce the impact of issues while at sea. Service support delivered through AR greatly extended the ability of onboard technicians to address failures they would have otherwise lacked the experience to diagnose and complete.

Our experts could identify issues from thousands of miles away and provide their maintenance crews with instructions to solve problems, thereby eliminating the need for re-routes, port stops and all the associated costs.

Connected tech

ABB'S AR solutions are available as downloadable apps from Google and Apple stores and use augmented reality to overlay the instructions on real equipment to expertly assist customers quickly and efficiently.

Interactive troubleshooting using step-by-step tutorials can be accessed by customers 24/7 for fast and easily accessible guidance through the different steps of key procedures. Facilitating remotely guided repairs and replacement of critical components takes this a step further, because in addition to using live on-screen annotations and digital overlays in the engineer's

field of vision, it also allows taking pictures, as well as audio and video sharing capability and guidance via live text chat.

What's more, the new Service Assist mobile app provides information even faster and more efficiently with the additional support of ABB-e, ABB's virtual assistant, which finds augmented reality immersive guides, books appointments for either on-site or remote services, and pulls together ABB and non-ABB documentation in a single digital location The speed of resolution is tangible and therefore vital in minimizing potentially highly disruptive and costly downtime.

Value added AR solutions should integrate multiple data sources and collaboration tools into the same augmented environment, so that teams can collaborate much more effectively, regardless of their location, and get virtually hands on. In our view, as we continue to navigate the challenges of a constantly changing energy landscape, the adoption of AR will only continue to accelerate. Why? Because providing operators with the visual information needed to fix problems and issues is a winning formula for achieving ongoing improved efficiency, quicker and safer resolutions, plus enhancing asset life and performance.

Quite simply, it empowers end users and boosts positive and proactive interaction, so that for both parties service support is conducted in a faster, more optimized and sustainable way.

Augmented Reality is very much a smart business reality!

AFL at DCW:
Advanced fibre network solutions

Data Centre World, ExCel London, March 6-7

AFL welcomes you to Stand D825 for the largest, thought-provoking event in the data centre calendar. Prepare for two jam-packed days of in-depth talks, mind-expanding expos, and unparalleled networking opportunities.

THIS YEAR, you're in for a treat. Not only can you expect to feast your senses on 100+ hours of interactive presentations and captivating panel debates—while celebrity spotting in the Data Centre World Keynote Theatre—but a hidden bonus awaits; your ticket grants entry to five leading technology events taking place simultaneously on site (Big Data & AI World, Cloud Expo Europe, Cloud & Cyber Security Expo, DevOps Live, and of course, Data Centre World). With around 15,000 industry professionals and end users gathered to take part in countless game-changing technology conversations, all under one roof, you won't want to miss the many insightful success stories from all corners of the industry (tip: business cards very much at the ready).

At the top of the bill is headline guest speaker Professor Brian Cox, who shares the stage with leading names from global brands heavily invested in the future of fibre, including Dr Carolyn Mercer (Chief Technologist, NASA), Zoe Kleinman (Technology Editor, BBC), and John Kundert (Chief Product and Technology Officer, Financial Times). Take our word for it, the epicentre of cutting-edge infrastructure opportunities comes but once a year – and you don't want to miss this one. Make sure you block out your diary and book your ticket well in advance. Tickets are free, but not unlimited (register at datacentreworld.com). Oh, and there's one more reason to book your ticket and dust off your favourite lanyard, and you can find it at Stand D825.

Stand D825 – where hyperscale happens

Join AFL's impressive portfolio of industry experts, on hand to help guide companies seeking innovative hyperscale data centre solutions in 2024 and beyond. Expect lively discussions, engaging demonstrations, and above all, the opportunity to network with AFL, currently celebrating 40 years of delivering premium, high-performance products and market-leading data centre services focused on meeting and exceeding customer expectations.

With decades of experience in the field, the AFL team is here to help kickstart your next hyperscale data centre project. From analysing your existing data centre environment to uncover innovative improvements across white space, sustainability, and productivity, to blueprinting customised, future-proofed infrastructures to suit your upcoming needs, AFL can assist you at every stage of planning, deployment, and aftercare. Here are some of the topics that the knowledgeable AFL team will be on hand to discuss in person:

AI networks

The expert AFL team brings many years of valuable expertise in high-performance networking. We deeply understand the pivotal role networks play in enabling AI at scale. Our purpose-built, end-to-end solutions help resolve even the toughest emerging challenges as your AI/ML workloads evolve in complexity. Whether you require high-throughput transport of petabytes across continents or real-time edge inference across millions of devices, our advanced cable routing platforms, high-density cabling, global manufacturing, and dedicated teams enable optimized network solutions at scale.

White space

Our goal is to provide your data centre with tailored white space solutions, ensuring your network remains manageable, flexible, and scalable. White space solutions from AFL ensure your network remains manageable, flexible, and scalable.

From upgrading your existing data centre infrastructure to architecting and implementing new deployments, our white space solutions can help futureproof your deployments. We offer fibre assemblies and modular housings to suit your individual network needs—speak to us at Stand D825 for more information about network planning, installation, and testing in any size data centre.

Data centre interconnect (DCI)

Our cutting-edge, scalable DCI solutions link multiple data centres to help you create integrated, shared-resource networks. Explore the advanced DCI solutions you need to help balance workloads, boost productivity, and accelerate deployment times while reducing overheads. With small-diameter, ultra-high fibre count cables, we fit more fibres into your existing duct space, preventing costly civil works. Choose between connectorized or spliced solutions (each crafted with state-of-the-art technology, including our innovative Wrapping Tube Cable with SpiderWeb Ribbon®) to ensure optimized network connectivity.

Unlock next-level modular connectivity

AFL offers tailored modular connectivity solutions designed to quickly launch new deployments or seamlessly enhance the performance of your existing fibre network. Whether you require high-density, tray-based, shallow-depth, or legacy application solutions, we have the answers you need for optimal fibre performance.

ASCEND® – Introducing ASCEND®, the high-speed, high-density, high-performance fibre connectivity platform for data centres. ASCEND® housings support both incremental growth and full-scale deployments in an easy-to-use, scalable platform.

U-series – Meet bandwidth requirements now and in the future with the fast, flexible U-Series fibre connectivity platform. The U-Series enables optimal network performance with high-capacity housings, easy-to-install cassettes, and high-performance assemblies.

H-series – Designed with optimal colocation data centre performance in mind. Future-proof your colocation fibre network, speed up deployments, lower installation costs, and reduce network complexity, all while increasing revenue by leaving more space for customers.

Understand and overcome hyperscale challenges AFL deeply understands the complex and varied challenges facing modern hyperscale data centres. From dependable access to advanced technologies to reliable global supply chains positively impacting your global and local availability, AFL can help you adapt and scale to the demands of your unique workloads. Whatever your next step, speak to us first. Visit Stand D825 to explore how we can assist you in 2024 and beyond.

Best practices for hiring

How should a business set about recruiting for senior positions? And what are the pitfalls to avoid to ensure securing the right person for the job?

BY PETER HANNAFORD, SENIOR PARTNER OF PORTMAN PARTNERS

THE PAST DECADE has seen data center capacity grow to unprecedented levels, particularly in the last few years, due to the surge in internet usage and low-interest-rate financing for major capital projects. The increasing use of artificial intelligence means that an almost unimaginable increase in internet growth is now predicted over the next five years and beyond. The largest internet companies are now engaged in an accelerating race to secure data center capacity in strategic geographies around the world – and are falling over themselves to be the winner in the game. Billions of dollars are being invested each week in growing the global digital infrastructure. The stakes have never been higher. We all know that the cost of poor leadership is staggering for businesses, so the pressure is on for those firms, their boards and their investors to find the best leaders to fill critical roles. It's certainly not an easy task, especially when there could be a large pool of potential candidates to consider.

Assessing talent, however, is much more complicated than simply looking for a candidate that ticks all the typical boxes, and often the best people are overlooked because they don't meet a narrow criteria.

So how should a business set about recruiting for senior positions? And what are the pitfalls to avoid to ensure securing the right person for the job?

Partner with your search firm

If you are working with an external search firm to find the best candidates that you can, (and obviously we'd recommend that approach!) then you should view the firm as a project partner, temporarily extending your own team and working together in full transparency and collaboration to work towards a common goal. There should be no territorial behaviour. All candidates, whether identified by the in-house recruiter or referred by a member of the existing staff should be considered alike on a level playing field. A senior executive at one of the hyperscalers urges both in-house and external search partners to think of the relationship as a market trade. External search partners offer market, competitor and candidate intelligence in trade for the in-house recruiter who offers the currency of hiring manager relationships, project management know how, and key messages from the employer brand. If both trade well, we all walk away with a successful search outcome.

The CV only tells you what people have done, not what they are capable of doing

Ask most people what they would cite as the worst hiring decision ever and they'd probably quote poor John Sculley's appointment at Apple. He'd enjoyed great success with PepsiCo but failed at Apple. Compare him to Jony Ive. He worked for a small design agency in London designing, among other things, toothbrushes, microwave ovens and toilets. He was hired by Apple and went on to design the iPod, iPhone and iPad – and was subsequently knighted. What did Steve Jobs see in Jony that wasn't in his CV? His HR team must have thought he'd gone mad. When it comes to picking a leader, one of the biggest mistakes is assuming that the right person needs to have had a particular job title or have specific academic qualifications. This may be true in certain cases (for example, technical roles where the requisite skills are only acquired with experience). However, a candidate may have all the skills needed to succeed in your leadership role but may simply not have had a chance to put them into practice - yet. Intellect, motivation, values, and behaviours are difficult to change. Experience is the one thing that constantly changes. Look beyond the CV.

Embrace Diversity

Diversity and inclusion are often lumped together as a single matter for consideration. Diversity is who organisations are, but inclusion is what they do to make diversity work. While many organisations have been busy ticking the diversity box by extending

opportunities to traditionally under-represented groups (in the digital infrastructure sector, diversity is often couched in gender terms) the underpinning rationale for this is often missed. While a data center business may cultivate some kudos by hiring more women, for example, getting more women into traditionally male dominated roles is not the end goal. Real value is generated from cognitive diversity – the heterogeneity of knowledge and intellectual perspectives that come from different life experiences and socio-cognitive variables like gender, cultural background, age, nationality, or even occupational expertise.

So, forget the box ticking often associated with surface level DE&I initiatives and get thinking about how systems, processes and culture can support diversity of thinking. And remember:

- Diversity reflects the real world
- Diversity means looking at the same idea in a different light
- Diversity brings disruptive ideas and innovation
- Your clients and customers are diverse.

Seek culture contribution, not culture fit

Research undertaken at Stanford University showed that, in a poll, 100% of participants put culture-fit above skills and intellect when reviewing candidates. And a study of 200 start-ups found those who put culture-fit first were least likely to fail and more likely to IPO. This is understandable as all had the same mindset, the same values and were uniquely motivated. However, the study also found that after going public the businesses grew at a slower rate than those who didn't rate culture-fit so highly. They were thinking as a group and had no diversity of thought. So culture-fit can be advantageous in some scenarios but not in all. Embracing diversity helps but ask yourself what's missing from your culture and look for candidates who can help improve it.

Have a clear, timely interview process

As stated earlier, if the candidate you're presented with is good then he or she will likely be targeted by other firms competing for their services.
Time is of the essence. Six or seven (or more) separate unstructured interviews are not only a burden on the candidate, but they are also ineffective with the same questions often being asked and the candidate having to repeat what they said in an earlier round. If the candidate is excited and motivated at first interview, they can, understandably, be somewhat demotivated by the seventh.

We've had candidates discounted in later stage citing "a tiredness or disinterest" as the reason for rejection. If you need multiple people involved in the decision making, then it is surely better to arrange interviews in a panel format. This is easy to organise using Teams or Zoom, or in the boardroom at later stages.

Better questions

When interviewing the candidate, avoid questions such as "why do you want to join us". Remember that you/we have approached the candidate. They didn't apply for the job. During the interview process you are still trying to get them to want to join.
A better question would be "what do you think you could do for us".

By all means test that they've done their homework; that they know exactly what your firm does and what the current challenges might be.

First impressions are important. What's your gut feeling when you spoke with or met this person for the first time? The way they looked, acted, spoke, and, importantly, listened. Were you impressed? As the saying goes, you only get one chance to make a first impression. What was your first impression? Having said that, don't decide that "this is the one" too soon. Keep the first impression but continue to ask probing questions.

Likewise, don't be unduly influenced by academic qualifications. Remember that Steve Jobs, Mark Zuckerberg, Michael Dell, and Bill Gates had no college degrees. Someone with industry experience is, of course, good but it's not predictive of success. If the person is locked into a fixed mindset and approach, industry experience can be a drawback. Old habits die hard, and you may have a big job to re-train the person to your company's way of doing things.

Make a shortlist

Ideally no more than three. Then make sure you check references.
It's likely that the candidate will still be employed, and you need to respect the confidentiality of the situation. Clearly you should never contact anyone without the candidate's express permission. If you are working with a specialist industry search consultant, they would be able to get discrete feedback from their own network. It's likely that the candidate will be known to them anyhow.

Some form of psychrometric test may be useful at this stage. Psychometric tests are used by employers to assess intelligence, abilities, potential and personality. Portman uses the GC Index. This is not a psychometric test but an organimetric, which measures and describes five different ways in which people are inclined to make an impact and contribution.

Finally

When you think you have found the right person, arrange for them to meet the board or their peers in a social environment. Perhaps a dinner. People act differently when relaxed and undoubtedly this person will be required to attend formal and informal events on behalf of your firm at some time. It also gives you the opportunity for this person to shine in the company of their peers.

Why legacy modernisation is essential to reduce energy costs

As the demand for digital infrastructure continues to increase, new-build data centres and a large estate of legacy IT facilities will be vital in underpinning the ongoing demand for digital infrastructure. The question is, how can we best meet today's sustainability challenges, alongside the equally important imperative of reducing energy costs?

By Andy Connor, EMEA Channel Director, Subzero Engineering

IN AN IDEAL WORLD, it would be nice to imagine that any data centre over a certain age would be quietly taken offline and retired – many being inefficient and ill-equipped to deal with today's increasingly high-density IT demands. However, in the same way, that the switch to electric vehicles allows for a long transition period away from fossil fuel-powered transport, to cushion the impact and, importantly, the cost of such change, the reality of the digital infrastructure sector is that legacy data centres will be required for the foreseeable future. All such facilities are capable of being replaced in a highly ambitious five-to-ten-year period, but the investment required would be prohibitive. Closing down older data centres without replacing them would mean a massive drop in digital infrastructure capacity, at a time when the demand has never been so high.

Establishing the exact global data centre population is not easy. In terms of what might be described as the number of data centres owned and/or operated by professional data centre providers, the US International Trade Commission estimates, as of May 2021, a total of some 8,000 data centres globally[1]. However, when considering the hundreds of thousands of enterprise-owned data centres globally – everything from a cupboard with a few servers, right up to large-scale buildings - then the extent of the legacy problem is truly revealed. Let's not forget that for every greenfield data centre, there's at least one older facility upon which one may rely.

Why does all of this matter?

Today, data centre PUEs have come down from an average of 2.5 in 2007 to 1.55 in 2022[2], but this current figure is still a long way from the ambition of 1.0. If the average is 1.55, that means that there's a substantial number of data centres that are not even close to that figure as of now, the impact of which comes at a significant financial and environmental cost.

Energy costs for the industry have risen massively since January 2021[3]. While the rate of the increase varies from region to region, many countries have witnessed a more than doubling of the cost of electricity. The silver lining to this energy crisis might be the accelerated development of renewable resources, for longer-term cost and environmental benefits, the reality is that, for many businesses, the price of power has become unsustainable.

In the data centre sector, where much of the cost of running a facility is the power bill, the impact has been significant. Smarter operators will pass power costs directly on to customers, which on the one hand, provides a haven from volatile energy costs, but on the other, it risks customers moving their business to cheaper or more energy-efficient colos,

as the overall operating costs will be cheaper
As for enterprise data centre facilities, the rise in energy prices across Europe is equally, if not more impactful, than it is for the colo and cloud operators. After all, enterprises are unlikely to be able to negotiate the long-term price agreements that have, to some extent, cushioned the energy crisis blow for large-scale operators.

For legacy data centre operators, the energy crisis is even worse. They are not only paying hugely increased power bills but are paying an additional premium through the inefficient use of power. Legacy data centres, for example, waste more than 60% of cooling energy[4], meaning higher power consumption, and a much higher energy bill than is necessary. Thanks to inefficient data centre design, and a lack of truly sealed, air-tight aisles of racks and cabinets, the cooling solution deployed in many data centres is having to work for much longer and harder than it would need to do in an optimised environment.

Containing the problem

The good news is that legacy data centres have a roadmap which can take them a long way on the journey towards minimising their carbon footprint and improving energy efficiency – something essential when considered as part of a wider sustainability programme. The first step is to analyse and understand what's going on in the data centre today. Without such insight, any attempted improvements will be little more than educated guesses, with no real idea of their likely impact.

CFD analysis provides a comprehensive map of how a data centre operates in terms of how effective the existing cooling technology is. It highlights key inefficiencies, such as hot spots, providing an overall picture of where hot air leaks in and/or out of the IT racks, cabinets, and aisles, demonstrating the inefficiency of the existing cooling solution. Armed with this picture of current legacy data centre performance, it is then possible to implement a programme of relatively simple, but extremely effective, measures to bring about a significant improvement. Most notably, the implementation of a comprehensive containment strategy brings with it significant data centre operational efficiencies. These efficiencies translate into significant financial and environmental benefits, saving as much as 30% of the data centre energy bill and offering a reduced carbon impact.

The emphasis is on 'comprehensive'. For example, there is no point in building a containment system around the IT cabinets, if the power cable ingress ducts are not fully sealed as well. Similarly, if a cabinet is not fully populated with hardware, the containment solution will not be effective unless the gaps are closed with blanking panels.

Properly designed and installed, containment systems can deliver huge operational and energy benefits to the data centre. Set alongside other measures, such as an optimised data centre configuration, increased use of free cooling where possible and/or increased operating temperatures, as allowed for in the ASHRAE standards, containment systems can improve legacy data centre performance dramatically.

Finally, comes the ongoing monitoring and maintenance requirement. Data centres are dynamic environments. The contents of racks and cabinets, hence the power supply required, will change over time. People will go in and out of the aisles, carrying out various tasks and, inadvertently or otherwise, impact the space. The performance of the IT hardware may also change over time, likely generating more heat.

The climate outside the data centre will impact the environment inside. In short, there are many ongoing variables in terms of data centre performance, and, at the most basic level, the IT load could vary dramatically on a daily basis.

Making modernisation work

Measure, contain and monitor is the mantra for an effective legacy data centre modernisation programme. Whatever the size of your data centre, understanding its current operational performance gives you a baseline from which to plan any improvements. Any such improvements will need at their core a commitment to a comprehensive containment solution. And ongoing monitoring and maintenance will ensure that the newly upgraded and optimised legacy data centre continues to run as energy efficiently as possible.

Optimising performance from an energy efficiency standpoint brings with it much-needed financial savings at a time of energy price volatility. Further, by using less energy, not only does a data centre cost less to run, but also enhances its sustainability credentials. Properly modernised and optimised, legacy data centres can continue to support the ongoing digital revolution and be seen as part of a sustainable, green future, and will no longer be seen as the industry's environmental Achilles heel.

FURTHER READING

➤ 1 https://www.usitc.gov/publications/332/executive_briefings/ebot_data_centers_around_the_world.pdf
➤ 2 https://www.statista.com/statistics/1229367/data-center-average-annual-pue-worldwide/?kw=&crmtag=adwords&gclid=CjwKCAjwjYKjBhB5EiwAiFdSfmIx0gH5seYY-3gxVaA7kti0zq1qiRhSuSeu2xk5ogn4Xs4Zo6t8xBoCAkIQAvD_BwE
➤ 3 https://www.consilium.europa.eu/en/infographics/energy-prices-2021/#:~:text=Graph%20showing%20energy%20prices%20in,went%20from%20105.4%20to%20172.3
➤ 4 https://www.datacenterknowledge.com/archives/2012/11/15/benefits-of-data-center-containment

UPS evolution and why sustainability is the next frontier

As an industry, we now need to look seriously at sustainability and how it can be genuinely achieved within datacentres and other critical organisations.

BY DAVID BOND, CHAIRMAN, CENTIEL UK

In the beginning...

TO UNDERSTAND the future, often we must look to the past. Modern humans evolved from primates, in various stages, over several hundred thousand years with each evolutionary stage having significant advantages over the previous stage. The uninterruptible power supply (UPS) industry has undergone a similar evolutionary process, but in a rather shorter time frame.

In the beginning (let's say 40 years ago), there were stand-alone UPS systems. These were huge monolithic units the size of a double wardrobe with a typical 60kVA (48kW) UPS weighing around 600kg.

When increased power availability was required (let's say 35 years ago) because of increasingly critical loads, parallel redundant UPS systems evolved. Back then, the only way to achieve parallel redundancy was to place two or more stand-alone UPS next to each other and install communication cables between each UPS.

When increased operating efficiency was required (let's say 30 years ago) because of increasing "environmentally friendly" awareness and the rising cost of energy, transformer-less UPS systems evolved. This evolution increased UPS operating efficiency by around 5%, decreased their size and weight, reduced their cost and enabled the next stage of evolution – modular UPS.

Around 20 years ago the first true modular UPS were developed, heralding the arrival of increased system availability, scalability, flexibility and maintainability. It is the development of these true modular UPS that we shall focus on now.

Where are we now?

Power protection systems must be available to provide clean, guaranteed power to their critical load every second of every day, so maximizing system availability must be the overriding objective of any installation. To achieve maximum availability, all single points of failure need to be removed.

Centiel's 4th Generation true modular UPS system, CumulusPowerTM, provides a significant improvement over previous designs. Each UPS module in the system contains all the power elements of a UPS – rectifier, inverter, static switch, display - and critically - all control and monitoring circuitry so there are no single points of failure.

Availability for Centiel's CumulusPowerTM UPS is further improved by the incorporation of Centiel's Distributed Active Redundant Architecture (DARA) and their Triple Mode communications bus. Duplication and redundancy of UPS components must also apply to the means of communication between UPS modules. The earliest, and simplest, communications bus is a single cable. However, if this breaks or becomes disconnected, the entire UPS system is compromised. For this reason, the ring circuit was introduced so if the communication cable was cut in one direction the signals could communicate the other way around the ring. The introduction of Centiel's Triple Mode communications bus increases system availability even further.

As the name suggests, Centiel's Triple Mode communications bus provides three separate paths of communication between UPS modules and frames made up of three separate ring circuits and three separate brains communicating with three other separate brains. It's the belt, braces and buttons approach where single points of failure are completely removed.

It is the combination of no single points of failure redundancy, DARA, Triple Mode bus and Centiel's safe/hot swap technologies that gives Centiel's 4th generation true modular UPS systems their class leading 'nine-nines' or 99.9999999% availability.

Scalability and flexibility are also important when installing a system to ensure the continual 'right sizing' of the UPS. A system which is too small will be overloaded, compromising reliability and availability, whilst one which is too large will be inefficient, will waste energy and will be unnecessarily costly to run. CumulusPowerTM has a

flat efficiency curve and provides greater than 97% efficiency even at low loads.

The next evolutionary stage

As we can see, the prevailing need has driven important developments in UPS technology over the years. The need facing us all now, and hence in UPS development, is Sustainability. Datacentres, for example, are an essential part of modern life but burn massive amounts of energy and use huge volumes of water. The earth is getting warmer and we are beginning to realise that some resources are finite. We must all work towards developing more sustainable solutions.

Centiel's development team has long been at the forefront of solving prevailing technology challenges. In the late 1980s I wrote a paper imagining the perfect UPS. I said it would be 100% efficient (we are now close to 98%), it would offer 100% availability (we are now at 99.9999999%), it would present a perfect load to the mains supply (it is now unity power factor with input current total harmonic distortion - THDi - of <1%) and would last forever. The development team behind Centiel, who created the first transformerless UPS and the first three phase true modular UPS, have pretty much solved the first three challenges to create the (almost) perfect UPS, and for the past four years they have been innovating to develop a UPS with a design life to match the design life of a datacentre (typically 30 years) and can be recycled to make UPS more sustainable.

Earlier this year, Centiel launched its new UPS, StratusPowerTM, which shares all the benefits of its award-winning three phase, true modular UPS CumulusPowerTM - including "9 nines" (99.9999999%) availability to effectively eliminate system downtime but now also offers class leading 97.6% on-line efficiency to minimise running costs and high quality, long-life components to improve sustainability.

Higher quality components cost more but this is just tens of pounds per module, compared to thousands to replace an entire system with a shorter design life. Combined with Centiel's approach as trusted advisors, it means they can help organisations take steps to move away from a "throw away" culture with a genuinely sustainable offering whilst helping to reduce Total Cost of Ownership (TCO) at the same time.

And finally...

Like all Centiel's UPS, StratusPowerTM is manufactured at its factory in Switzerland. However, uniquely, it includes even higher quality components so instead of replacing filter capacitors and cooling fans every four years, they now need replacing every 15 years, or just once during StratusPower's entire 30 year design life. Furthermore, at end of life, StratusPowerTM can also be almost 100% recycled.

Like Centiel's whole range of UPS solutions, UPS cabinets are designed with scalability and flexibility in mind, and future load changes are easily accommodated by adding or removing UPS modules as required. An organisation will never outgrow a well specified StratusPowerTM UPS and it can be constantly rightsized to ensure it always operates at the optimal point in its efficiency curve.

Centiel's UPS solutions also work with Lithium Iron Phosphate (LiFePO4) batteries which tolerate higher ambient temperatures, further reducing the need for cooling, and will also only need to be replaced once in a 30 year design life.

As an industry, we now need to look seriously at sustainability and how it can be genuinely achieved within datacentres and other critical organisations. It is the next step in UPS evolution and StratusPowerTM is a significant move towards this goal.

Centiel's experienced team is always available to advise on the most appropriate options for organisations looking to improve their approach to sustainability and reduce their carbon footprint while maintaining the highest level of availability for critical power protection. Centiel can help organisations take steps to move towards genuine sustainability.

The team at Centiel has been at leading-edge of UPS development and evolution over many decades. Centiel's goal is clear: to achieve the ultimate availability of power for its client base. Its quality technology, backed-up with comprehensive maintenance contracts ensure Centiel's clients' have the very best power protection at all times, now and in the future.

For further information please see:
www.centiel.co.uk

Data center performance measurement that adds up to improved efficiency

Modern data center operations extend beyond merely ensuring IT functions run smoothly. Today, data center operators face the growing challenge of optimizing their facilities for maximum efficiency.

BY MARK FENTON, SENIOR PRODUCT MARKETING MANAGER AT CADENCE

WHILE THE TERM “efficiency” is a recurring theme in the data center industry, its interpretation can vary. Is efficiency about conserving energy use, cost reduction, water conservation, carbon emissions or something else entirely?

Each person you ask may give you a different answer. Still, when speaking to senior management, all these factors are likely to contribute to their definition of efficiency.

By employing physics-based analysis, data center managers can truly understand how to tackle all the different factors that are in play in order to ensure their facility runs as efficiently as possible.

Managing the demands of a data center

Delivering higher processing throughput, whilst meeting the various demands across a facility – such as constrained space and sustainability goals – is a common challenge data center operators face for a number of reasons.

Firstly, in a modern data center, there is often a growing frequency of high-density IT equipment deployment requests that puts huge pressure on the facility systems. This high-density equipment often necessitates additional power and cooling resources that older facilities, particularly legacy ones, aren’t designed to accommodate.

In addition to this challenge, some data centers lack the technology to predict the impact that operational changes can have elsewhere. For example, a manager could deploy high-density IT hardware and be unaware of how this will affect the current cooling and power system provisioning.

This lack of visibility can increase risk that can only be comfortably resolved by over-provisioning once the new equipment is in place. This would have a negative impact on the data center’s sustainability posture as it would mean a rise in the carbon

emissions produced. With the introduction of regulations, such as the Corporate Sustainability Reporting Directive in Europe, which necessitate stringent tracking and reporting of power and cooling use, facility managers can't afford to have their carbon footprints surge unnecessarily.

Faced with balancing new sustainability requirements with demands already in play, how can managers ensure their data centers run at the best performance level possible? One route is through measurement, which can help identify how different demands interact with one another and have the most significant effects.

Scientific data center performance measurement
Measuring data center performance with The Green Grid's Performance Indicator (PI) plus computational fluid dynamics (CFD) simulation is a proven approach to enhance efficiency in a facility and troubleshoot any issues.

The Green Grid's PI is a powerful tool that empowers managers to establish the data centers' ability to meet IT cooling needs and be efficient so they can improve performance. It does this by delivering visibility into how different factors – including energy use, cooling effectiveness, and IT equipment performance – impact each other.

More specifically, the Green Grid's PI can help weigh-up effective cooling during normal operation (IT Thermal Conformance) and adequate cooling during failure or maintenance within design parameters (IT Thermal Resilience), without compromising efficiency (PUE ratio).

Meanwhile, CFD technology simulates cooling systems so data center managers can virtually test operational changes and measure the impacts before they're made in the physical facility. Combined, this approach offers managers a holistic view of how any given change will impact the data centers' performance. This means they can better plan, test, and validate changes, taking into account both risk and efficiency, before any change is made in the real data center.

Improving data center efficiency

Uncovering small improvements that can be made in a data center, can massively enhance a site's overall efficiency and, therefore, the data center's performance. The first step in discovering what these adjustments are is measurement, which quantifiable approaches, like

The Green Grid's PI and CFD simulations, make simpler and more effective. Executed well, performance assessments empower managers to balance the conflicting demands within their facility, ultimately making their job of handling operations simpler and the data center more efficient.

Sustainability must roll forward faster

At last, almost everyone agrees that the world's environment is in crisis. Governments, industries, corporations, down to individuals are now taking steps towards better environmental stewardship. We must all take responsibility for our actions and help reduce the CO_2 emissions we create and work to achieve Net Zero.

BY MICHAEL AKINLA, BUSINESS MANAGER, NORTHERN EUROPE, AND IRELAND, PANDUIT

SUPPLIERS, customers, investors, and employees all understand that this is not an easy path, but it is an essential one. We have all witnessed, over the past decade, how the data centre industry has grown from an obscure market to a key economic engine for every developed and developing country. The data centre industry continues to grow and with it the need to reduce the negative environmental impact generated, from the construction phase, through fit out, to resource utilisation, such as, water and energy and back through the supply ecosystem.

At Panduit, we concentrate on delivering the physical ITE infrastructure to data centres, enterprise and industry and environmental stewardship is part of a holistic approach to the company's business philosophy. We have realised that customers recognise that sustainability is a key driver for future growth, and they want to partner with companies that are committed to sustainable business practices.

We, like many organizations have implemented a broad strategy of greening our supply, production processes and offices, and this is not a late to market effect of customer pressure. These changes were initiated over a decade ago, demonstrated by our Costa Rica manufacturing facility which has been carbon neutral certified since 2014. The plant was recognised as a pioneer in CO_2 reduction at the 2019 PreCOP25 International Climate Conference.

A key goal is to recycle thousands of metric tons of manufacturing waste per year and reduce Panduit's carbon footprint by 20% by 2025. A further manufacturing goal is to reduce energy intensity (kWH/Earned Hours) by 10% by 2025, across our facilities.

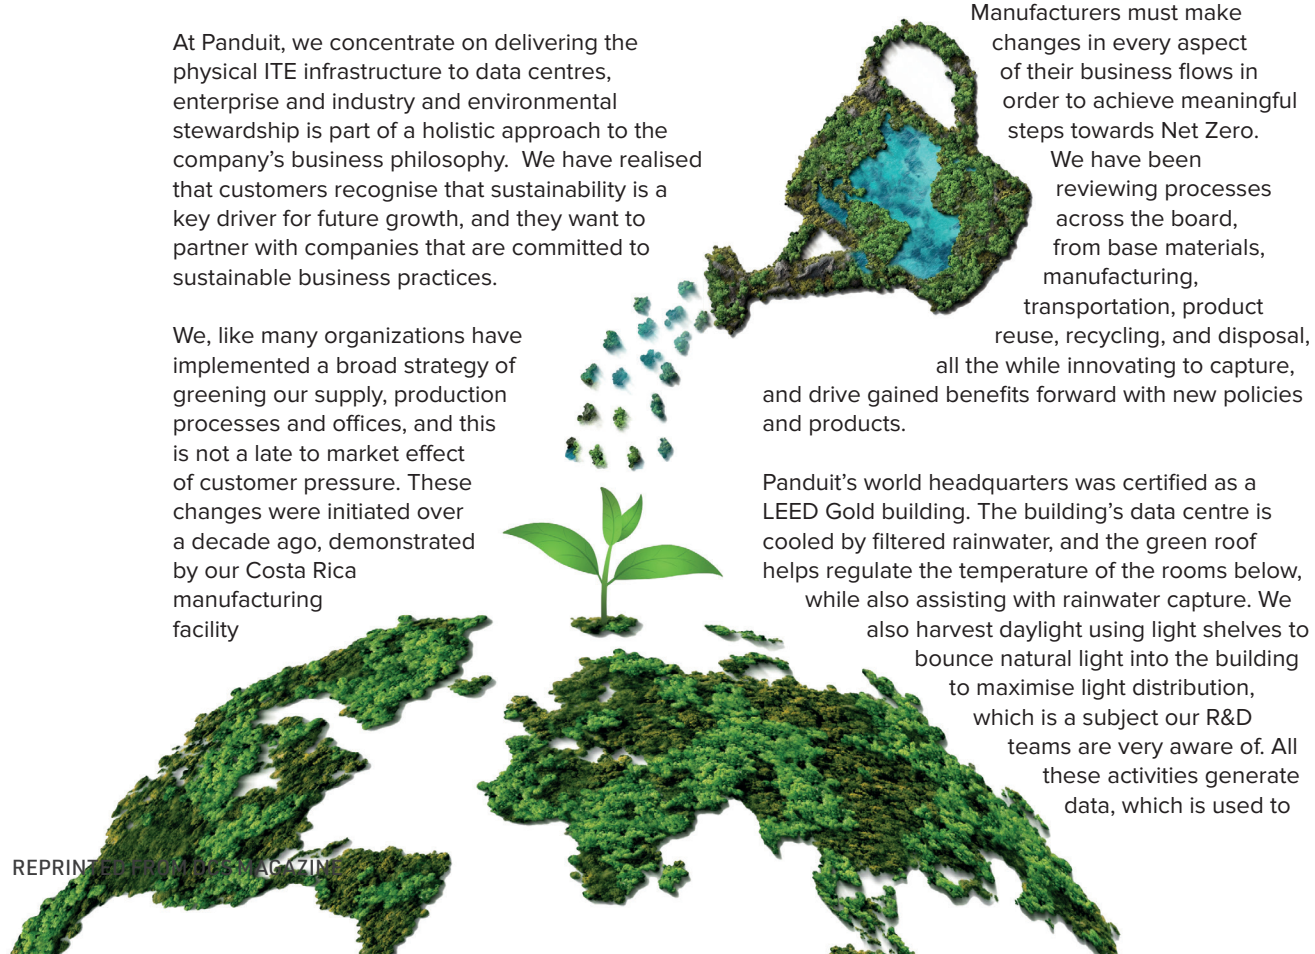

Manufacturers must make changes in every aspect of their business flows in order to achieve meaningful steps towards Net Zero. We have been reviewing processes across the board, from base materials, manufacturing, transportation, product reuse, recycling, and disposal, all the while innovating to capture, and drive gained benefits forward with new policies and products.

Panduit's world headquarters was certified as a LEED Gold building. The building's data centre is cooled by filtered rainwater, and the green roof helps regulate the temperature of the rooms below, while also assisting with rainwater capture. We also harvest daylight using light shelves to bounce natural light into the building to maximise light distribution, which is a subject our R&D teams are very aware of. All these activities generate data, which is used to

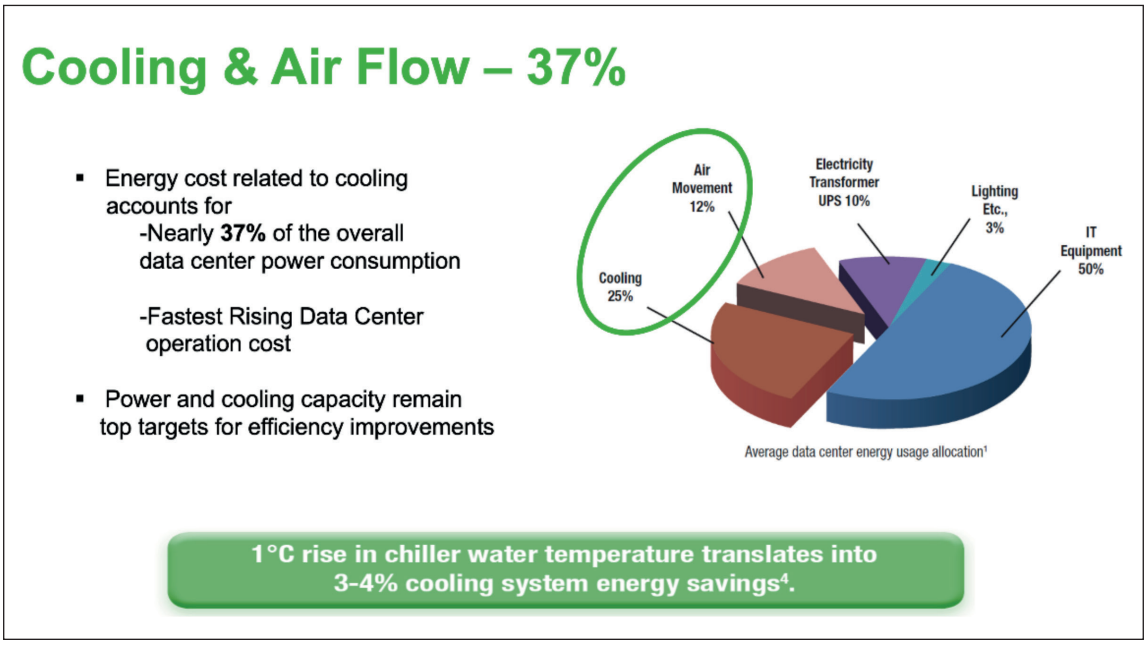

➤ Data Center Energy Allocation (Panduit)

benchmark, evaluate, and improve the operation of the building's systems in order to conserve energy, improve occupant comfort, maximise productivity and enhance operations.

Recycling for all

Organizations must be strident in material sourcing. We are RoHS and REACH compliant and have implemented recyclable material packaging. We work with our suppliers to upgrade packaging increasing renewable wood and cardboard products as well as recyclable plastic materials.

We, like many organisations, measure, analyse and report the environmental and health impacts of copper cable and connector materials as part of the Environmental Product Declaration (EPD) and Health Product Declaration (HPD) process.

Current investigations include recycling excess fluorinated ethylene propylene (FEP) back into base material and the reuse of copper waste generated from the manufacturing process to achieve 100% waste recycling.

E-Waste is a major challenge according to the Global E-Waste Monitor project. 53.6 Megatons (Mt) of e-waste was generated around the globe in 2019, of this only 9.3 Mt, or 17.4% was disposed of or recycled properly.

Panduit's R&D Labs have been investigating product efficiency for almost 60-years. We are helping data centre customers be more sustainable, whether on-premises or standalone sites, by sharing our knowledge and research concerning the amount of energy wasted within infrastructure environment. Up to 40% of data centre energy is used in cooling and lighting, therefore, increasing energy efficiency is important in improving sustainability and reducing carbon emissions.

White is right

A number of years ago, our labs undertook research into the benefits of white cabinets compared to black. It found that white cabinets provided a Light Reflective Value of 80% verses black cabinets that offered 5%, this equates to a light energy saving of up to 30%, (1.5-3% of data centre energy), and requiring fewer light fittings in the area. Panduit's research also found that the reflective light allowed easier working in and around the cabinets for installation and MACs.

An even bigger opportunity to increase energy efficiency is around the servers. ASHRAE TC9.9 has encouraged operators to raise temperatures in technology spaces and introduce environments such as Hot/Cold Aisle Containment Systems that eliminate hot air recirculation. Improving cold air direction into cabinets intake fans reduces inlet air temperature and improves cooling and equipment reliability. Ensuring cabinets are air sealed ensures empty rack spaces and gaps around cable inlets are covered to reduce air leakage by up to 20% and

Panduit's R&D Labs have been investigating product efficiency for almost 60-years. We are helping data centre customers be more sustainable, whether on-premises or standalone sites, by sharing our knowledge and research concerning the amount of energy wasted within infrastructure environment

Lighting – 3%

- White Cabinets
  - Light Reflective Value (LVR) 80% vs.5% for Black
    - 25 to 30% light energy savings
  - Increase Visibility inside the cabinet
    - Ambient light sufficient for most work

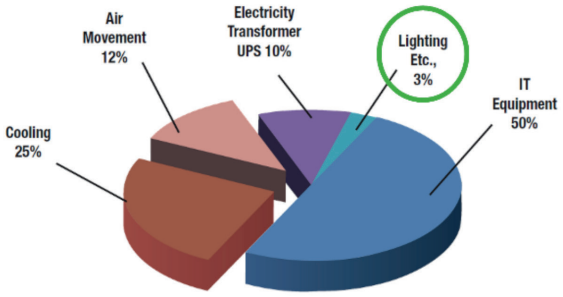

Average data center energy usage allocation[1]

preventing hot air contamination. These measures can reduce energy use by a staggering 40%, a demonstrable sustainability win.

Today, monitoring devices and DCIM systems ensure that data in and around cabinets is collected and reportable to ensure cabinets operate at maximum efficiency, which is good for the data centre, good for customers and good for the environment. Panduit continues to review product materials and sourcing strategies. Vari-Matrix HD 28AWG patch cords offer the industry's smallest diameter Cat 6A cable, which uses less material in its manufacture, and maintains the quality and capability of its characteristics. This also allows more cables per bundle around the site and in distribution more reels per pallet, increasing shipping efficiency by reducing the number of traffic journeys.

Companies looking to evaluate potential suppliers should ask for ESG policies and evidence they contribute to internal efficiencies and drive lower energy use, reduction in base materials usage and reduced packaging. Also key, is what internal practices are suppliers taking to drive internal efficiencies to improve sustainability drives, such as EPD/HPD certification on euroclass rated cables which help with LEED and BREEAM and demonstrate a cradle to grave measurement approach.

TMC

world of transformers

EMPOWERING THE DATA CENTRE WORLD WITH TMC TECHNOLOGY

TMC STANDS AT THE FOREFRONT OF THE TRANSFRMERS WORLD, OFFERING INNOVATIVE SOLUTIONS TO CATER TO THE NEEDS OF DATA CENTRES WORLDWIDE

STAND D329

TMC Transformers SpA - Viale dell'Industria 65, 21052 Busto Arsizio (Italy) - tmctransformers.com

Space and power: what's really limiting data centre capacity

Perceptions of limitations in data centres may not be what they first appear, as demand drives new examination

By Markus Gerber, Senior Business Development Manager, nVent Schroff

THE WORLD has seen increasing demand for digital services in recent decades. This demand has only grown since the pandemic when they were not just a boon but a lifeline for many during times of lockdown and isolation.

As we moved beyond the pandemic, work practices were also changed forever, as more and more sought to work remotely, whether from home or different geographies. This combined with developing business models and digital transformation, has seen demand grow even more, as well as new requirements that have seen the likes of edge computing proliferate.

All of this has driven growth in data centres but has also increased pressure to meet demand. These pressures are seeing space, density, and power come under the spotlight as limitations in some cases.

Scale of growth

To get an idea of the scale of growth in digital services in recent decades, data volume is a key indicator. According to Statista, since 2010, the volume of data created, captured, and consumed has grown from 2 zetabytes to 97 zetabytes in 2022, with the figure for 2025 expected to be 181 zetabytes. Despite this near exponential growth, according to the International Energy Agency, energy demand since 2010 has only gone from 194 terawatt hours (TWh) to just over 200 TWh in 2022.

These two contrasting figures show the extraordinary strides that have been made in energy efficiency in computing since then, especially when it comes to pure processing power. With Moore's Law in effect for the period, the benefits are clear. Now though, there are concerns from no less a figure than the CEO of Nvidia, Jensen Huang, that the Moore's Law effect may be coming to an end. While this is disputed, there can be little doubt that processors are likely to become evermore powerful, while producing more heat in the process. To meet that demand towards 2025, and beyond, it is likely that data centre limitations will be encountered, with space being chief among them.

Space and power

Space was often seen as one of the chief limitations for data centres. The space was described in terms of the ability to power equipment in a given unit of measure, such as watts per square foot or square meter. This was a useful rule of thumb for specifications and facility design. Architects would plan cooling and power according to such measures. Under this approach, the data centre has progressively been becoming hotter, using more power, to provide for increasing levels of processing. In an air-cooled data centre this required more and more air pumped through, meaning for every watt drawn, less and less was used for compute.

As a result, data centres in the nineties and early 2000s became less and less efficient in terms of how effective that power was in being used to provide data processing. As the chip power kept going up through a number of different technological developments, and with evermore demand for performance, data centre operators found themselves demanding more and more cooling volume and flow, until they hit barriers in cost, complexity and management. Many experienced a threshold where they simply could not just pump a room full of air to be able to cool those chips, making it increasingly unfeasible for much of what is already deployed.

Equipment management

Management too became an issue. Often as a data centre evolved, equipment was upgraded or altered, moved around or replaced due to failure. Gaps, spaces, and expansions often meant that even carefully implemented methodologies such as hot aisle/cold aisle systems, were left working poorly, as guidelines for airflow management were often ignored in the need for expediency and demand.

This could add to the impression of space limitations when a new project or service was contemplated, when in fact a facility, if properly managed, could take more before reaching the inevitable limit of pumped air cooling.

What is clear from this is that while good management and design are key to ensuring that physical space is not a limitation in meeting demand for digital services from data centres, air cooling is and will be increasingly so in the future. As other architectures also emerge, such as edge computing, new cooling solutions will be required if service demand, physical space, energy efficiency and sustainability needs are to be met.

To meet the emerging demands for digital services in the foreseeable future, data centre operators will need to consider hybrids of air, liquid and direct-to-chip cooling, taking advantage of the specific characteristics of each to appropriately and proportionately provide the kind of cooling that enables density to be deployed reliably and economically.

Inefficient medium

There is a clear reality when it comes to cooling: the closer heat can be captured from where it is produced, the more efficient the process.
Allied to this is the fact that air is a very inefficient medium. A water-based fluid, or another dielectric liquid, is a much more efficient medium to capture and transmit heat.

Even with the likes of hot and cold aisle layouts, with rear door and in-row coolers, blanking plates and efficient cabling, air is still inefficient. While these measures are likely to remain part of the mix for many operators for years to come, other methods must be considered. Liquid cooling solutions can offer greater capability to accommodate equipment density than air cooling, Heat captured through liquid cooling can be more efficiently removed from the immediate environs of the equipment, and brought to potential reuse opportunities, without a state change.

Developments available now in liquid and direct to chip cooling can not only meet today's density demands, relieving physical space limitations, they can also offer a critical upgrade path to allow data centre operators to move towards more efficient methods. This will be crucial as budgets also come under pressure amid the ongoing inflation trends and continuing global uncertainty.

Strengths and purpose

With these new cooling techniques and systems, it is not a one size fits all approach. Each technique and system has its particular strengths and characteristics that must be taken into account ensure the right performance is delivered per requirement. In-rack, in-row, or direct precision liquid cooling all offer differing applications and benefits to achieve and overall density and performance goal.

All the while, ensuring efficiency that contributes to sustainability targets. Data centre operators must be supported in their design and operational objectives by a trusted technology partner that not only has in depth knowledge but also a broad portfolio of solutions to meet each need. Understanding where better managed air cooling can remain, liquid cooling can be adopted and direct to chip cooling leveraged, is key to getting current needs under control, while building a path to future capability.

Improvements and a path forward

By properly examining real or perceived data centre space limitations, data centre operators can determine how best to tackle their density needs. More efficient, precise and controllable cooling solutions will be a key part of that effort. With efficiency as a central strand of sustainability efforts, hybrid systems of air, liquid and direct cooling techniques can build a path to greater effectiveness in data centre cooling that relieves space pressures, while meeting demand and providing a strong base for future growth.

Is liquid spray cooling the future for data centres

Maintaining the temperature of critical infrastructure within Data Centres through the use of air-cooling technology has dominated the industry for many years, whilst the demand in data consumption has increased year-on-year and in turn has increased the amount of electricity consumed.

ELECTRICITY CONSUMPTION for IT and Cooling systems accounts for approximately 86 percent of the total energy consumption within a Data Centre, while cooling alone can account for up to 40 percent.

Furthermore, in recent years we have seen server technology evolving at a fast pace, resulting in an increase in heat density. Typically, heat density increases by an average of 1kW per rack every two years and we are now starting to reach the point where the effectiveness of air-cooling technology, in some cases, is restricted due to the air's heat transfer coefficient, which limits its ability to remove heat from today's modern chips. Dealing with this issue has been at the forefront of industry debate. The solution in overcoming these limitations is Liquid Cooling.

Some Data Centres are already adopting Liquid Cooling, but the specific technology that has so far been predominately implemented is immersion cooling. Although it is an efficient method in removing heat from CPUs, the technology still has some reliance on mechanical cooling, which makes it less sustainable, and limits the ability of the Data Centre in moving closer to net zero. Combined with the growing pressures of the global energy crisis and repurposing waste heat from the Data Centre, what could be the solution?

Established in 1995, Airsys is a market-leading cooling solutions provider, who think globally, but act locally. We deliver innovative, high-efficiency, precision control thermal solutions for the built environment. With over 25 years' experience, combined with multiple manufacturing facilities and offices globally, Airsys are able to deliver sustainable solutions for critical environments such as Data Centres and Telecoms environments.

A focus on understanding the customers' needs, combined with our technological expertise, has allowed Airsys to develop a liquid spray cooling solution, called the LiquidRack™ which addresses the current limitations of existing technology and moves Data Centres closer to achieving their sustainability goals.

LiquidRack ™is a liquid spray cooling solution designed for multiple types of digital data infrastructures, such as cloud service, telecommunication facilities and more. Differing from immersion cooling, our pioneering approach keeps the dielectric fluid moving, spraying directly onto the CPU. Adopting a liquid spray cooling approach significantly increases the heat transfer coefficient,

in comparison to immersion cooling. Increasing the heat transfer coefficient, allows the dielectric fluid being sprayed on the CPU to be elevated up to a temperature of 65oC, without comprising performance of the CPU. This provides two major benefits. Firstly, the elevated fluid temperatures eliminate the need for chillers, and they can be replaced by dry coolers in environments with ambient temperatures up to 50oC. This means that there are potential free-cooling applications in geographic locations such as the Middle East, for which free-cooling has been traditionally hard to achieve. Secondly, the elevated water temperatures allow for heat recovery.

Heat recovery means that the recovered energy from the Data Centre can be recycled for multiple applications, such as district heating and industrial scale greenhouses to name but a couple. The Data Centre now becomes an energy producer, helping to achieve Net Zero emission goals. LiquidRack™ can offer a seamless connection for heat recovery to district heating systems, due to the high running temperatures, which is another advantage to the solution.

The largest obstacle for the adoption of liquid cooling technology has always been the significant expense involved in transitioning away from air-cooled solutions in existing Data Centre environments. The project complexities and capital expenditure were previously considered to be high-risk. LiquidRack™ changes this, now the transition is seamless. LiquidRack's™ vertical design, size, and lower operational weight, when compared to immersion technology, allow for easier integration into existing Data Centres. Each LiquidRack™ consists of two drawers, with each drawer, for our standard design, able to accommodate up to 6 x 2U servers, therefore each system can hold up to 12 x 2U servers. LiquidRack™ has been designed to provide a cooling capacity of up to 50kW. Each server can be slid out and locked into position for hassle-free inspection and maintenance.

The system is fully adaptable to different server brands, different server layouts and even a mixture of CPUs and GPUs within the same drawer. Designed as a decoupled system, each drawer contains two low powered pumps (N+1) and a highly efficient heat exchanger. The LiquidRack™ design offers an unbeatable fantastic cooling capacity when compared to its modest footprint.

As the LiquidRack™ directly cools the CPU via spray technology, not relying on immersing the server, the dielectric fluid needed is typically reduced by 80% when compared to immersion technology. This provides CAPEX advantages in terms of dielectric fluid procurement and reduces structural issues in relation to weight.

By spraying directly onto the CPU, the LiquidRack™ provides a constant flow and eliminates uneven

fluid flow that can occur using other methods of liquid cooling technology, which can result in server reliability issues. Uneven flow can occur when there is a mixture of different servers immersed or if one or more servers are removed for maintenance, resulting in the dielectric fluid following through the path of least resistance, thus starving the CPU of the flow needed to remove heat. Therefore, LiquidRack™ offers a more reliable solution to preserve servers and maintain Data Centre uptime.

As data consumption becomes ever more important in our day-to-day lives both personally and professionally, global Data Centre electrical consumption is predicted to reach 4% of the total global electrical consumption by 2030. A combination of the world's reliance on gas and oil, and the surging costs associated with them, turbulent energy prices affecting operating costs and budgetary planning, and the need to move towards a more sustainable future, are driving change in the industry.

The Data Centre community needs to quickly adapt and Airsys passionately believe they have a game changing technology that provides a low CAPEX and low OPEX solution that can turn a Data Centre into an energy producer, whilst achieving free cooling anywhere in the world. Liquid Cooling is the next step for the Data Centre industry to achieve their aims of sustainability and Net Zero, and liquid spray cooling is the pioneering next step in the liquid cooling journey.

www.air-sys.uk

As data consumption becomes ever more important in our day-to-day lives both personally and professionally, global Data Centre electrical consumption is predicted to reach 4% of the total global electrical consumption by 2030

The liquid future of data centre cooling

With rising demand, and equipment densities, air as a cooling medium is reaching its limits. Developments in hybrid and liquid cooling will allow providers to rise to the challenge sustainably.

By Markus Gerber, Senior Business Development Manager, nVent Schroff

THE DATA INFRASTRUCTURE industry is facing a number of challenges in today's digital world. Demand for data services is growing at a phenomenal rate and yet, there has never been a greater pressure, or duty, to deliver those services as efficiently and cleanly as possible.

As every area of operation comes under greater scrutiny to meet these demands, one area in particular, cooling, has come into sharp focus. It is an area not only ripe for innovation, but where significant progress has been made that shows a way forward for a greener future.

According to some estimates, the number of internet users worldwide has more than doubled since 2010, while internet traffic has increased some

20-fold. Furthermore, as technologies emerge that are predicted to be the foundation of future digital economies, such as streaming, cloud gaming, blockchain, machine learning and virtual reality, demand for digital services will rise not only in volume, but also sophistication and distribution. Increasingly, the deployment of edge computing, bringing compute power closer to where it is required and where data is generated, will see demand for smaller, quieter, remotely managed infrastructure. This one area alone is expected to grow at a compound annual growth rate (CAGR) of 16% to 2026 to a market of more than $11 billion, according to GlobalData, a research, consulting and events business.

This level of development brings significant challenges for energy consumption, efficiency, and architecture. The IEA already estimates that data centres and data transmission networks are responsible for nearly 1% of energy-related greenhouse gas (GHG) emissions. While it acknowledges that since 2010, emissions have grown modestly despite rapidly growing demand, through energy efficiency improvements, renewable energy purchases by information and communications technology (ICT) companies and broader decarbonisation of electricity grids, it also warns that to align with the Net Zero by 2050 Scenario, emissions must halve by 2030.

This is a significant technical challenge. Firstly, in the last several decades of ICT advancement, Moore's Law has been an ever-present effect. It states that compute power would more or less double, with costs halving, every two years or so. As transistor densities become more difficult to increase as they get into the single nanometre scale, no less a figure than the CEO of NVidia has asserted that Moore's law is effectively dead. This means that in the short term to meet demand, more equipment and infrastructure will have to be deployed, in greater density. Added to this, are the recent developments from both Intel and AMD, where their high-end data centre-aimed processors will work in the 350-400W range, further exacerbating energy demand.

All changes will impact upon cooling infrastructure and cost

In this scenario of increasing demand, higher densities, larger deployments, and greater individual energy demand, cooling capacity must be ramped up too.

Air as a cooling medium was already reaching its limits, being as it is, difficult to manage, imprecise, and somewhat chaotic. As rack systems become more demanding, often mixing both CPU and GPU-based equipment, individual rack demands are approaching or exceeding 30W each. Air-based systems, at large scale, also tend to demand a very high level of water consumption, for which the industry has also received criticism in the current environment. One estimate equated the water usage of a mid-sized data centre as equivalent to three average-sized hospitals.

Liquid cooling takes many forms, but the three primary techniques currently are direct-to-chip, rear door heat exchangers, and immersion cooling

Liquid cooling technologies have developed as means to meet the demands of both volume and density needed for tomorrow's data services. Studies with different liquid cooling techniques have established that they can be anything from 50 to 1,000 more efficient than air cooling.

Liquid cooling takes many forms, but the three primary techniques currently are direct-to-chip, rear door heat exchangers, and immersion cooling. Direct to chip (DtC), or direct to plate, cooling is where a metal plate sits on the chip or component, and allows liquid to circulate within enclosed chambers carrying heat away. This is a highly effective technique, that is precise and easily controlled. It is often used with specialist applications, such as high performance compute (HPC) environments.

Rear door heat exchangers, as the name suggests, are close-coupled indirect systems that circulate liquid through embedded coils to remove server heat before exhausting into the room. They have the advantage of keeping the entire room at the inlet air temperature, making hot and cold aisle cabinet configurations and air containment designs redundant, as the exhaust air cools to inlet temperature and can recirculate back to the servers. The most efficient units are passive in

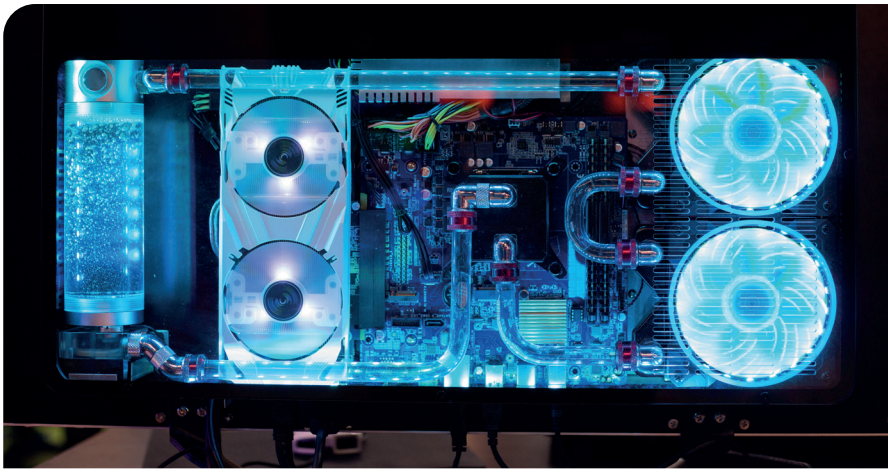

nature, meaning server fans move the air necessary. They are currently regarded as limited to 20 kW to 32 kW of heat removal, though units incorporating supplemental fans can handle higher loads in the 60 kW maximum range.

Immersion technology employs a dielectric fluid that submerges equipment and carries away heat from direct contact. Whilst for many, liquid immersion cooling immediately conjures up the image of a bath brim full of servers and dielectric, precision liquid immersion cooling operates at rack chassis-level with servers and fluid in a sealed container. This enables operators to immerse standard servers with certain minor modifications such as fan removal, as well as sealed spinning disk drives. Solid-state equipment generally does not require modification. A distinct advantage of the precision liquid cooling approach is that full immersion provides liquid thermal density, absorbing heat for several minutes after a power failure without the need for back-up pumps. Liquid capacity equivalent to 42U of rack space can remove up to 100 kilowatts (kW) of heat in most climate ranges, using outdoor heat exchanger or condenser water, allowing the employment of free cooling.

Cundall's liquid cooling findings

According to a study by engineering consultants Cundall[1], liquid-cooling technology consistently outperforms conventional air-cooling, in terms of both PUE and water usage effectiveness (WUE). This, says the report, is principally due to the much higher operating temperature of the facility water system (FWS), compared to the cooling mediums used for the air-cooled solutions. In all air-cooled cases, considerable energy and water is consumed to arrive at a supply air condition that falls within the required thermal envelope. The need for this is avoided with liquid-cooling, it states. Even in tropical climates, the operating temperature of the FWS is high enough for the hybrid coolers to operate in economiser free cooling mode for much of the time, and under peak ambient conditions, sufficient capacity can be maintained by reverting to 'wet' evaporative cooling mode.

A further consistent benefit, the report adds, is the reduction in rack-count and data hall area that can be achieved through higher rack power density. There were consistent benefits found, in terms of energy efficiency and consumption, water usage and space reduction, in multiple liquid cooling scenarios, from hybrid to full immersion, as well as OpEx and CapEx benefits.

In hyperscale, co-location, and edge computing scenarios, Cundall found the total cost of cooling information technology equipment (ITE) per KW consumed in liquid versus the base case of current air cooling technology, varied from 13-21% less. In terms of emissions, Cundall states PUE and Total Power Usage Effectiveness (TUE) are lower for the liquid-cooling options in all tested scenarios. Expressing the reduction in terms of kg CO2 per kW of ITE power per year, results saw more than 6% for co-location, rising to almost 40% for edge computing scenarios.

What does the immediate future hold in terms of liquid cooling?

Combinations of liquid and air cooling techniques, in hybrid implementations, will be vital in providing a transition, especially for legacy instances, to the kind of efficiency and emission-conscious cooling needs of current and future facilities. Though immersion techniques offer the greatest effect, hybrid cooling offers an improvement over air alone, with OpEx, performance and management advantages.

Even as the data infrastructure industry institutes initiatives to better understand, manage and report sustainability efforts, such as the Climate Neutral Data Centre Pact, the Open Compute Project, and 24/7 Carbon-free Energy Compact, more can and must be done to make every aspect of implementation and operation sustainable. Developments in liquid cooling technologies are a significant step forward that will enable operators and service providers to meet demand, while ensuring that sustainability obligations and goals can be met. Initially hybrid solutions will facilitate legacy operators to make the transition to more efficient and effective systems, while more advanced technologies will ensure new facilities more efficient, even as capacity is built out to meet rising demand.

By working collaboratively with the broad spectrum of vendors and service providers, cooling technology providers can ensure that requirements are met, enabling the digital economy to develop to the benefit of all, while contributing towards a liveable future.

Endnote

1 "Desktop Study Report - Liquid and Air-Cooling Compared," Cundall, March 2021

INTELLIGENT COOLING FOR CRITICAL ENVIRONMENTS

AIRSYS

WWW.AIR-SYS.UK
ENQUIRIES@AIR-SYS.UK

DCS AWARDS 2024

CELEBRATING 14 YEARS OF SUCCESS

The DCS Awards: 39 Categories across 5 Themes

23 MAY
Leonardo Royal Hotel
London St Pauls

KEY DATES:

23 FEBRUARY: ENTRIES WILL CLOSE AT 17:30 BST

29 MARCH: SHORTLIST ANNOUNCED
The list of finalists for the 2024 DCS Awards will be made public on the website

01 APRIL: VOTING OPEN
Voting open to subscribers of the Digitalisation World and Data Centre Solutions

03 MAY: VOTING CLOSES

23 MAY: AWARDS CEREMONY

Category Sponsors

H2H COMMUNICATIONS

oryx align

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

VOTE NOW: https://dcsawards.com/nominate

Supported by

dca The data centre trade association

Empowering Data Centre Expansion: The Impact of TMC Technology

In the dynamic landscape of the digital age, where information is the currency and data centres serve as the backbone of the digital economy, the role of technology in powering and supporting these centres cannot be overstated.

TMC TRANSFORMERS, a leading multinational manufacturer, specializing in medium and low voltage dry-type transformers, stands at the forefront of this technological revolution, offering innovative solutions for data centre developments across the globe.

TMC Transformers boasts a global presence and is known for its commitment to excellence in design, manufacturing and service. Leveraging the latest developments in cast resin and VPI (Vacuum Pressure Impregnation) technologies, the company delivers a comprehensive range of products that are tailored to meet the specific requirements of various industries, including data centres.

Data centres are the heartbeat of the digital economy, critical assets for storing, processing and distributing vast amounts of data generated by millions of interconnected devices. With the exponential growth of data-driven technologies, the demand for efficient and reliable data centre infrastructure has never been higher. TMC Transformers recognizes the critical role played by data centres in driving economic growth and innovation and is dedicated to supporting their expansion through cutting-edge solutions.

At the centre of TMC's data centre offering is a commitment to efficiency, reliability, and performance optimization. The company's extensive portfolio of medium voltage (MV) and low voltage (LV) dry-type transformers is designed to address the unique challenges faced by data centre operators. This includes non-linear loads with high harmonic content, transient voltages, space constraints, noise control, and leakage current management.

TMC Transformers' solutions for data centres are characterized by their strength, environmental sustainability and maintenance-free operation.

TMC's transformers offer unparalleled performance and reliability in the most demanding environments. TMC Transformers' innovative H-ART technology is at the pinnacle of transformer insulation systems, enhancing performance and extending asset lifetime expectancy of even further

With features such as the climatic classification E4-C5-F1, advanced glass resin insulation systems, low loss innovations that are compliant with EU 548:2014 EcoDesign Tier 2 standards, K-Factor design to K-13 specifications, electrostatic screens and enclosed solutions, up to and including IP67.

TMC's transformers offer unparalleled performance and reliability in the most demanding environments. TMC Transformers' innovative H-ART technology is at the pinnacle of transformer insulation systems, enhancing performance and extending asset lifetime expectancy of even further.

This advanced technology ensures optimal operational efficiency and reliability, helping data centre operators meet the growing demands of their customers while minimizing downtime and operational costs. As the demand for data centres continues to increase, TMC Transformers remains committed to driving innovation and excellence to support the industry. The company's participation in leading industry events such as Data Centre World London, emphasizes its dedication to staying at the forefront of technological developments and fostering collaboration with key stakeholders in the data centre ecosystem.

In conclusion, TMC Transformers plays a vital role in supporting the growth of data centers worldwide through its innovative and reliable solutions.

By leveraging the latest technologies and engineering expertise, TMC empowers data center operators to meet the evolving demands of the digital economy while ensuring maximum efficiency, reliability and sustainability.

As data centres continue to evolve and expand, TMC Transformers stands ready to conquer the challenges of tomorrow's data-driven world.

www.tmctransformers.com

EcoStruxure™
Innovation At Every Level

Stay CONNECTED

To your critical Infrastructure with EcoStruxure™ Data Centre solutions.

Keeping pace with the growth of demanding compute needs and digital transformation makes speed, resiliency, and connectivity critical for data centres.

Expand rapidly to meet your ever-changing needs with EcoStruxure data centre solution that scales to size and speed — in the Cloud and at the Edge.

Discover how EcoStruxure Data Centre solution delivers efficiency, performance and predictability. Meet us at booth D730 at Data Centre World London to find out more!

se.com/ecostruxure-datacenter

Life Is On | Schneider Electric