



INNOVISION

INSIGHTS + PERSPECTIVES



DELIVERING INTELLIGENT,
SUSTAINABLE INFRASTRUCTURE

IN ASSOCIATION WITH



ISSUE II 2021
<https://datacentre.solutions>

ROUNDTABLE

WITH ERIK SALO AND ANDY PALMER

10:00 – 11:00 AM (+1 GMT)

REGISTER NOW

<https://digitalisationworld.com/s/seagate>

In this exclusive roundtable, tech celeb Erik Salo and Seagate enterprise solutions specialist Andy Palmer discuss the latest intelligent storage solutions to optimize data centers.

They go into depth on Seagate Corvault, Salo's invention, and reveal the key trigger that led to it. Afterwards, they will be addressing questions.



Erik Salo
Sr. Director
Worldwide Systems
Distribution



Andy Palmer
UK Channel Lead
Enterprise Data
Solutions

A New Category of Intelligent Storage

Exos CORVAULT is a high-performance, self-healing block storage system that delivers multi-petabyte capacity, five-nines availability, and hyperscale efficiencies for data center and macro edge environments.



Editor's View

By Phil Alsop



Delivering intelligent, sustainable infrastructure

COP26 is nearly upon us. Opinion might be divided as to whether this event is rightly or not billed as something of an environmental last chance saloon; but we would at least appear to have travelled far enough along the green road whereby few climate change sceptics remain. However, taking everyone on the next stage of the journey will not be easy. Most, if not all, organisations have been more than happy to sign up to what might be called the 'green window-dressing' of the past few years: activities such as offsetting carbon emissions, signing up for renewable energy, becoming involved with recycling initiatives. These have been the relatively easy environmental wins. What's now required is nothing short of massive changes across all aspects of business (and private lives).

Working life during the pandemic has given some indication of what the future might hold in terms of eliminating, or at least minimising, vast quantities of business-related travel. The virtual world, made possible by data centres, keeping us all in touch when physical meetings were not possible. However, not everyone seems to be convinced that the office commute should come to a permanent end. And we've not even begun to explain to all those folk who were so desperate to go on an overseas holiday just as soon as lockdown was over that leisure travel, with the attendant carbon footprint, might just not be acceptable sometime in the not so distant future...

The data centre and wider IT industry have a paradoxical role to play in a sustainable future. They are major power consumers;



but they also provide many of the solutions which can help organisations and individuals to reduce their carbon footprints in spectacular fashion. So, as we look into the future, the data centre industry faces a twin challenge: to reduce its power consumption and overall environmental impact significantly whilst also underpinning the IT solutions which can help us all to become greener. King Solomon himself might have struggled with such an apparent paradox.

I believe that the contents of our InnoVision publication gives cause for optimism. Sustainability is very much front and centre for the data centre industry. Whether or not Net Zero is achievable and if so, in what sort of timeframe, we can rest assured that the data centre ecosystem is working towards such an objective with real purpose.

DCS **DATA CENTRE SOLUTIONS**

Editor Philip Alsop	+44 (0)7786 084559	philip.alsop@angelbc.com
Sales Manager Peter Davies	+44 (0)2476 718970	peter.davies@angelbc.com
Sales Executive Jessica Harrison	+44 (0)2476 718970	jessica.harrison@angelbc.com
Design & Production Manager Mitch Gaynor	+44 (0)1923 690214	mitch.gaynor@angelbc.com
Publisher Jackie Cannon	+44 (0)1923 690215	jackie.cannon@angelbc.com

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

Directors
Stephen Whitehurst – CEO
Scott Adams – Joint - Managing Director
Sukhi Bhadal – Joint - Managing Director

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



Datacentre Solutions is published 4 times a year on a controlled circulation basis in Europe, Middle East and Africa only. Subscription rates on request. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd. © Copyright 2021. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. ISSN 2396-9016 (Online)

DC BYTE HEADING

06 Looking ahead to 2022 and beyond

ENERGY OPTIMISATION

- 09** The Sustainability Paradox
- 11** Optimising Energy - How a Different Approach to Cooling Supply Temperature Can Lead to Increasing Efficiency
- 12** Energy Optimisation and Sustainability
- 14** Green Colocation - Pushing the Boundaries of Data Centre Sustainability
- 16** Sustainable Data Centres Need an Actionable Strategy
- 18** Meeting Demands for Data Centre Performance and Sustainability
- 20** Why Data is the Key to Sustainable Supply Chains
- 22** Data Centres as Partners of the Energy Transition
- 23** Arm-Based Servers can Significantly Reduce the Carbon Footprint of Data Centres
- 25** How Data Centres can Help Businesses be More Sustainable



HYBRID ARCHITECTURE

- 28** Why Cloud Native Postgres is the Natural Next Step in Databases
- 30** Putting the Future of Data Centres at the Edge
- 31** Distributed Computing and Kubernetes - Getting Data Everywhere
- 32** 5 Best Practices for Edge Infrastructure Management
- 33** Unlocking the Edge Opportunity: The Role of Regional Data Centres
- 34** On Hybrid Architecture

COLOCATION + HOSTING

- 36** Adopting a Hybrid Model
- 38** Data Centres are in Transition, What Choice Will Organisations Make?
- 39** Colocation in a COVID World
- 40** 8 Ways to Remotely Manage Your Colo Infrastructure
- 42** Is Colocation a Feasible Option for Your Business?
- 44** Colocation is Forever



06



14

AI + AUTOMATION

- 46 Dynamic Workload Provision Demands Automated and Adaptable Data Centre Power Provision
- 48 Automation via Integration: 3 Real-World Examples
- 50 Utilise AI or Risk Losing Out to the Competition
- 51 How AI and Automation Will Impact Data Centres in 2022

DESIGN + INFRASTRUCTURE

- 54 If You Can't Stand the Heat...
- 55 Unleash the Power of Flash with Fibre Channel
- 56 Is Object Storage the New Primary Storage?

56



60

PHYSICAL SECURITY + SAFETY

- 59 Seven Layers of Physical Security for Data Centres
- 60 Database Security is of Utmost Importance - and Open Source is the Answer

Looking ahead to 2022 and beyond



Despite, or perhaps because of, the challenges of the last 18 months, the data centre industry has been fortunate to experience continued growth. Take-up across EMEA rose 4% in Q1 2021 to 120MW with a 10% increase in new supply overall, totalling more than 180MW. In APAC total supply increased by just under 200MW, reflecting a similar pace to 2020 take up - recorded at just over 800MW - making total supply 5800MW across the region.

BY WILL KING, MANAGING DIRECTOR, DC BYTE

LOOKING AHEAD, DC Byte's data suggests that take-up rates will continue to rise, quickly absorbing 2020's live supply. This will potentially heighten demand in Tier II markets which reported little supply growth in early 2021. It also confirms the need to deliver on existing commitments to bring new data centre facilities online in 2022 and beyond.

Competition to source new sites to meet this demand will continue to increase exponentially and we expect to see the corresponding rise in land prices persist in many of the established markets for several reasons. Investors are searching for more profitable property investments away from traditional, but poorly performing, real estate assets such as the office market.

'Land grabs' are becoming more commonplace, as data centre developers strive to secure suitable sites which will protect their ability to develop new facilities either in the short or medium term.

The availability of competitively priced land and the continued trend for large data centre deployments in key locations has made it prohibitive for smaller investment firms to enter the market but private equity companies are continuing to invest in the sector,





even for larger sites. That said, although the core markets of Amsterdam, Frankfurt, London, Paris and Dublin maintained momentum, with London reporting 40MW of new supply in Q1 2021, the trend towards expansion outside of these markets, is expected to continue. Istanbul and Warsaw are already noted as markets to watch. Our analysis of private equity announcements also shows that Africa is expected to receive a high level of investment in the coming years. The establishment of new cloud regions and expansion of cloud operators into new territories to support their ambitious growth plans, will also drive demand for facilities in new and emerging markets.

Addressing this demand at a time when consolidation continues in the industry, and potential restrictions on development in some countries are put in place, will require even greater focus on strategic site selection when exploring opportunities beyond the established markets. It will also require efficient allocation of resources and skills needed to design, build and operate these facilities.

It's for this reason we expect the trend towards self-build to continue amongst hyperscalers as they elect to take ownership of the entire process, from design and specification through to ongoing operation. Until recently, we typically saw the highest level of self-build sites being announced across Dublin, the Nordics and the Netherlands. However, the hyperscalers are becoming more comfortable in developing across more geographies.

Build-to-suit will continue to grow as a model for the hyperscalers to the detriment of traditional wholesale colocation. Cloud companies will continue to deploy via data centre operators where access to resources in emerging markets could otherwise prove to be a barrier to entry, or where speed to market is an essential requirement.

We have seen a rise in 'power banking' over the last 12 months, with operators taking more power than they immediately need in order to accommodate future requirements. This is especially the case in larger markets, including London and Frankfurt, where public cloud operators have a presence and where there is concern that it may become harder to access power in the timeframes and specific locations needed.

DC Byte's research shows that as these challenges arise in EMEA, we expect them to be replicated across APAC, which is now seeing significant interest beyond its core markets. One fundamental difference between the two regions is the wider level of competition between the cloud operators in APAC giving clients greater freedom of choice.

Without question the data centre sector continues to move at a pace and is definitely one to watch. There is now a real need for investors, real estate experts and data centre operators to have access to real-time market intelligence to deliver in what will continue to be a highly profitable but extremely competitive sector.

ENERGY OPTIMISATION

Encouragingly, the industry as a whole has made great progress in addressing the issue of energy optimisation, with initiatives at everything from individual company level upwards. But there's more work to be done, as digital transformation places ever more demands on IT and data centre infrastructure resources.



IN ASSOCIATION WITH

Nick Ewing
Managing Director, EfficiencyIT
Efficiencyit.com



The sustainability paradox

IN RECENT YEARS data centres have come under increasing scrutiny for both their impact on the environment and the demands they place on energy infrastructure. An article in Science Magazine estimates that today data centres account for around 1% of worldwide electricity use. However, given that capacity requirements globally are continuing to advance at a rapid rate, energy demands are likely only to increase. A new report from Knight Frank, for example, states a 10% increase in new supply across EMEA, which totals 180MW of new supply.

There also remains the environmental challenges of recycling equipment that has reached its end of life, and the growing CO2 emissions from data centres at the edge. The question, therefore, has to be asked, can the world afford to sustain the number of data centres that are currently in place, and what changes are needed for those being built in the future?

Certainly, the 'sustainability' issue is one of increasing concern, and is being addressed by regulators, and legislators across the industry itself. The Climate Neutral Data Pact, for example, adopted by many hyperscale and colocation providers in Europe, commits its signatories to become carbon neutral by 2030, guiding new and legacy operators alike to attain greater operating efficiencies, and implement measurable sustainability targets.

Greater attention is also being paid to both optimising PUE (power usage effectiveness) ratings, and the need to implement energy efficiency or modernisation programs. Further still, the growing use of mandatory renewable energy, and a willingness to experiment with microgrids or reusable on-site energy storage is another positive contributing factor to sustainability. A paradox remains, however. Legacy data centres, those most



IN ASSOCIATION WITH

in need of improving their power consumption, efficiency and sustainability ratings, are among those least likely to procure the technologies to enable them to do so.

Regardless of the pronouncements of those concerned with Corporate Social Responsibility (CSR), where sustainability is a key concern for the operator and their suppliers, the harsh reality is that tender documents for modernisation programs rarely, if ever, make any mention of a need to meet sustainability criterion.

The fact is that sustainability ambitions at the C-level are a far cry from many procurement teams, and when this issue arises, it imparts the view that operators don't necessarily want the most efficient technologies; they want the cheapest.

Until sustainability becomes a quantifiable, achievable element of a Request for Quotation (RFQ) or Request for Procurement (RFP), it will remain, for the majority of legacy data centres, simply something for other people to worry about. New data centres can indeed be built for sustainability, but legacy sites remain the issue.

Changing this mindset will require a concerted effort throughout the supply chain, from technology vendors to operators, and

from the C-Level to procurement. Sustainability can only be achieved via the right design, embracing circular practices and by using the right technical infrastructure.

Changing the business case to focus on long-term strategic outcomes, such as reduced operating expenses, lower TCO, or decreasing carbon emissions is far more effective than simply shaving off infrastructure costs. Such an approach requires greater collaboration, communication and transparency.

Regrettably, sustainability is not free; it requires investment. There is a price to be paid for procuring better management software, more advanced UPS systems and by paying greater attention to inefficient mission-critical systems in need of modernisation.

There are many benefits to be gained in the long term, and a joined up, transparent and collaborative approach will not only help the industry thrive, it will significantly reduce its impact on the environment.

What's clear, however, is that while procurement are still incentivised to ignore sustainability in favour of the lowest cost-of-purchase, it will remain ambiguous, and change will be slow. Sustainability doesn't need to be a paradox, but the focus on ROI needs to change before it will become a reality.



DCS ONLINE ROUNDTABLE

BASED around a hot industry topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion.

MODERATED by an editor, this online event would include 3 speakers, with questions prepared and shared in advance.

THIS ONLINE EVENT would be publicised for 4 weeks pre and 4 weeks post through all our mediums and become a valuable educational asset for your company

Contact: Jackie.cannon@angelbc.com



Charlie Bater

Associate Director B&W Engineering

bandwengineering.com



Optimising Energy:

How a different approach to cooling supply temperature can lead to increasing efficiency

IMPROVING ENERGY EFFICIENCY is seen as a key metric for data centre operators due to a reduction in operational costs, environment impacts, and capital plant required to deliver the electrical demand. There are many opportunities for improving efficiency in energy demanding data centres, by both directly reducing the demand or indirectly by using otherwise wasted energy for other processes.

One direct measure, with large energy saving potential, is increasing supply temperature of the critical cooling infrastructure. Historically, reliability has been the major concern, with over cooling servers preferred to the risk of providing insufficient cooling.

In recent years cooling supply temperatures within the industry have steadily increased, with the 2011 revision of ASHRAE standard TC9.9 being one of the main catalysts for this change, with the recommended range changing to be between 18 and 27 degree Celsius (an increase of 2 degrees Celsius on the upper limit).

Benefits of increasing supply temperature

Increasing the supply temperature of the critical load, allows for the cooling medium (i.e., chilled water) temperature to be increased, which in turn allows for a higher number of hours of 'free-cooling', as opposed to full refrigerant cooling or partial free cooling.

For example, for a traditional chiller system:

1. Full refrigerant cooling utilises the compressor and refrigerant circuit within the chiller to provide cooling capacity at high ambient temperatures.
2. Full-free cooling mode bypasses the refrigerant circuit and provides the full cooling capacity via free-cooling coils within the chiller that operate at low ambient temperatures, which

can typically occur when the ambient temperature is 10 - 12 K below the chilled water leaving temperature.

3. Partial free-cooling is a mix of refrigerant mode & full free-cooling mode whereby some of the capacity is provided by free-cooling and some is provided by the refrigeration circuit. Partial free cooling typically occurs 4 – 5 K below the chilled water entering temperature down to the point of full free-cooling.

To give an example of the benefits, comparing a chilled water system with 20 °C Flow & 30 °C return against a system of 16 °C Flow / 24 °C Return for an average weather year in Frankfurt, Germany, yields the following results:

Cooling Type 20°C Flow / 30°C Return 16°C Flow / 24°C Return

Free Cooling Hours	4560	2860
Partial Free Cooling Hours	4090	5100
Full Refrigerant Cooling Hours	110	800

Trends and Predictions

Although the ASHRAE TC9.9 standards recommended range limits temperature to 27 degrees Celsius for all classes of servers (A1-A4), the allowable range permits supply temperatures between 32-45 degrees Celsius depending on the class of equipment.

Allowing the IT equipment to operate at the allowable range (32°C for A1) creates significant energy reduction compared to the recommended range of 27°C. However, to realise further energy reductions within data centres, it is likely and deemed imperative that tenants and server manufacturers will continue to improve the ambient temperatures their servers can operate at.



Stuart Crump
Director of Sales, Iceotope
iceotope.com

Energy optimisation and sustainability

INEFFICIENT AIR-COOLING SYSTEMS in data centres are not going away any time soon. The vast majority of live sites, and many data centres in the planning stages, still utilise similar systems with varying degrees of inefficiency. Most organisations planning for small, remote edge locations are still clinging to 'white space', tried and tested, chilled air and fan assisted compute that perpetuates energy inefficiency.

Cooling constraints continue to inhibit compute capacity, even as direct greenhouse (GHG) emissions regulation comes into play, and although there are more energy efficient alternatives available for data centres and edge applications.

A year of disruption and the global shutdown caused by Covid-19 has ceded to a year which will be known for Data Gravity, a phrase that describes how data location and generation attract more and more applications and services to it. Given that IDG indicates data volumes are growing at an average of 63 percent per month, and by 2025 over 463 exabytes of data will be created each day, thus illustrating the continued requirement for new data centres.

Disruption can be good and increasing numbers of legacy data centres are realising that while their air cooling satisfies their 2-3kW server racks, it cannot meet the requirements of High Performance Computing (HPC). However, reconfiguring technology halls to accommodate liquid cooled HPC pods allows air and liquid cooling to co-exist more efficiently and opens new business opportunities. It can also positively impact the sites PUE or ITUE, depending on how accurately you like to measure your ITE energy effectiveness.

As compute density continues to grow, so the heat generated by the ITE equipment increases and as the industry moves forward the removal of heat will become more regulated. Therefore, we need solutions that allow heat reuse, local energy partnerships and infrastructure in order to share or resell that commodity to the community. Liquid cooling removes heat at a much higher temperature than air cooling, around 50°C, and

technologies are becoming available to convert that heat into stored energy to be reused. Thus, creating a revenue stream for the data centre owners and reducing the overall OPEX.

Couple the data growth with Gartner's predicted 50:50 model, where fifty percent of data will be generated and utilised at the user location, therefore necessitating a more agile and adaptable mindset in the design of data centres whether large or small, remote or in town centres. Data centres can no longer be grouped together around high availability power nodes, such as West London.

Across the country, where power distribution is limited, the data centre industry must be more effective in targeting where power is utilised within edge facilities. This means compute capability must take priority, and ancillary activity, such as cooling, must support energy efficiency, not detract from it.

Many cooling techniques require compressor-based air cooling; the technology is understood within the industry, it's relatively reliable and resilient, however, it is by anyone's reckoning, inefficient. Continuing to deploy air-cooling based edge data centres is simply not sustainable.

In light of the latest Scope emissions protocol, our edge data centres must be installed as highly efficient users of energy. Eliminating air cooling chillers and fans can liberate around thirty percent of the energy used by the site, and without the need to pull the chilled air through the server, the rack fans can also be removed, which could save around ten percent of the server's power use.

Data gravity, coupled with the growth of data centres of all descriptions, has consequences for the whole industry's supply chain. Liquid cooling is not industry critical yet, however with legislators targeting sectors with high energy and land use, the market must demonstrate its understanding of the requirement to implement technology that will help optimise energy use and improve sustainability.

Bring digital to every person, home and organization
for a fully connected, intelligent world





Lee Myall
CEO, Kao Data
kaodata.com

Green colocation - Pushing the boundaries of data centre sustainability

LARGE OR SMALL, disaggregated or centralised, data centres continue to be fundamental to the digital economy. In many respects, the carbon impact of our sector has gained much attention within the last decade, yet sustainability and energy efficiency have long been at the forefront of the industry.

While average global PUE's remain flat, according to the Uptime Institute, many operators continue to push the boundaries of efficiency and innovation, something of paramount importance as power-hungry GPU-accelerated computing proliferates. Further, inward initiatives such as the Climate Neutral Data Centre Pact have driven greater collaboration and awareness on an international scale, thus ensuring that Net Zero targets become both measurable, effective, and most crucially achievable.

Across the industry, many hyperscale and colocation providers are changing long held doctrines around design, cooling and operational methodologies. All of which can make significant contributions to reduce the global impact of climate change. What's clear, however, is that every single organisation within our sector must continue to do more, evaluating and exploring new ways to drive sustainable operations for the good of the planet. Picking a certified green energy tariff is far from enough.

At Kao Data, we believe that sustainability begins before the design stage, and from inception it was one of our guiding principles. Data-driven insights have long been instrumental within digital infrastructure, and by utilising digital twin and CFD (computational fluid dynamics) technologies, we have been able to model the design and build cycle under numerous operating conditions. From high performance computing (HPC) and artificial intelligence (AI) to enterprise colocation, we have been able to analyse our design with remarkable accuracy, grounding deployment decisions in hard facts that support sustainability imperatives.

Cooling, for example, is pivotal for green colocation, and by implementing mechanical and refrigerant-free indirect evaporative cooling, we're able to reduce carbon dioxide, harmful fluorinated gas emissions, and monitor gaseous contaminants. Such an approach ensures that server and storage warranties are never put at risk, while offering an industry leading PUE of 1.2, even at partial loads.

Moreover, as chip and processor technologies evolve to offer greater performance benefits, so must the cooling technologies and infrastructure systems which support them. Many data centres remain in a quandary over how to support HPC and AI workloads without major structural work. The ability to accommodate OCP-Ready™, customised architectures with access to 'direct-to-chip' liquid cooling, therefore, provides





a strong and energy efficient foundation to support the most complex of machine learning and supercomputing workloads.

Furthermore, liquid cooling technologies provide a means to re-use waste heat, something that Kao Data hopes to explore in the future. In areas where the technology converges with local plans for sustainable developments, heat re-use could enable district heating networks to be deployed on a much wider scale. This is something that could be transformative for UK Net Zero ambitions, and within Smart City initiatives such as those taking place in Cambridge.

Finally, renewable energies are crucial, but today they come in many shapes and forms. Today Kao Data contracts 100% sustainable energy, yet data centres with ambitions of achieving Net Zero must remain vigilant, and be prepared to test ground-breaking technologies that result in carbon reduction. As I mentioned earlier, it's not enough to just pick the green tariff, as

an industry, we must conduct root and branch work to eradicate all fossil fuel usage, where possible.

Recently, we have eliminated the use of fossil fuels from our campus, using HVO (Hydrotreated Vegetable Oil) to power our backup generators. Created primarily from recycled cooking oil, HVO is a high-purity, advanced and synthetic, renewable fuel, which is 100% sustainable and non-toxic. It results in a 90% reduction in Greenhouse Gas (GHG) emissions and positively supports customer demands to reduce Scope 3 emissions.

Further, it shows how our industry can take another simple step toward reducing CO₂, while reusing its existing infrastructure. As we look forward, the pathway to carbon neutral data centres requires innovation, openness and an appetite for collaboration. At Kao Data, we believe that the more we push the boundaries of green colocation and share our findings with the sector, the more it will benefit customers and the surrounding environment.

As chip and processor technologies evolve to offer greater performance benefits, so must the cooling technologies and infrastructure systems which support them. Many data centres remain in a quandary over how to support HPC and AI workloads without major structural work



Marc Garner

VP, Secure Power Division, UK&I, Schneider Electric

se.com

Sustainable data centres need an actionable strategy

THROUGHOUT THE LAST YEAR digital transformation has grown at an unprecedented rate. Connectivity and data centre demands are accelerating, and our world is fast-becoming a digital domain. With business, economies and consumers growing ever dependent on mission-critical digital infrastructure, it's paramount that the data centres underpinning them become more sustainable.

In late 2019, the Harvard Business Review stated 99% of large company CEOs agree "sustainability issues are important to the future success of their businesses." Yet today sustainability has also become a key topic within our sector. Internet Giants like Apple, Facebook, and AWS have all made public commitments to sustainable operations, and Google, for example, is racing to become the first company to operate carbon free by 2030.



Many of the world's most influential operators are also embracing sustainable practices. Companies such as Microsoft, Interxion and Kao Data have all become signatories to the Climate Neutral Data Centre Pact, setting ambitious targets to become carbon neutral by 2030, while helping to influence policy and regulation at an international level.

Sustainability drivers

Over the last decade, technology and efficiency improvements have undoubtedly provided us with a crucial vehicle to reduce carbon emissions and energy losses, but those alone will not be enough to deliver the next leap in sustainable data centres. For many years data centres have been seen as large consumers of power, with various analysis suggesting they represent 1-2% of global electricity consumption. Indeed the Schneider Electric Digital Economy and Climate Impact Report states that IT sector electricity demands will grow by 50 percent by 2030, reaching 3,200TWh, equivalent to 5% Compound Annual Growth Rate (CAGR) during the next decade.

A report from Schneider Electric and 451 Research, however, found that 97% of colocation customers are demanding contractual commitments to sustainability, and of the +800 global operators surveyed, 55% were also taking some action, but it was neither strategic nor comprehensive, leaving a clear gap between customer expectations and providers' commitments.

What's clear is that as we look to build the sustainable data centres of the future, operators need an actionable approach that changes the dynamics of traditional deployment strategies. This includes embracing opportunities to integrate with the grid and becoming ingrained with the circular economy, but how will operators achieve this?

Defining a framework

At the first step, operators need to agree and define a bold,



actionable and integrated strategy; empowering design, procurement, facility and sustainability teams to collaborate and reach a common goal. At the second, they must utilise resource efficient data centre designs, which use fewer raw materials and offer an extended lifecycle through technology innovation. During the third step, operational efficiency becomes essential. Here owners and operators must utilise data to gain visibility of critical systems, track energy usage and benchmark their performance against said measurable targets.

At the fourth, renewable, or energy management strategies, encompassing renewable energy credits (RECs), on site, or distributed power generation via microgrids, or offsite generation via power purchase agreements (PPAs) is paramount to reach the next stage of their sustainability objectives.

At the fifth and final stage, it's critical we decarbonise the supply chain. Scope 3 emissions are the industry's biggest challenge and I believe that Vendors who embrace the circular economy, and are transparent about the sustainability impact of their products, can take one step towards helping operators understand their embodied carbon footprint.

For sustainable data centres to become a reality, it's crucial that actionable strategies take precedence over words. As an industry, no one company has all the answers, but it's clear to see we're making progress.

Over the last decade, technology and efficiency improvements have undoubtedly provided us with a crucial vehicle to reduce carbon emissions and energy losses, but those alone will not be enough to deliver the next leap in sustainable data centres. For many years data centres have been seen as large consumers of power, with various analysis suggesting they represent 1-2% of global electricity consumption



Andy Connor

Channel Director, EMEA, Subzero Engineering

subzeroeng.com

Meeting demands for data centre performance and sustainability

WHEN IT COMES to the deployment of energy efficiency programmes, there's a clear divide between modernised and legacy data centres.

On the one hand, hyperscalers, and indeed many colocation providers, are designing and building ultra-efficient digital infrastructure systems, 'embedded' with energy efficiency and sustainability best practices. All these operators have access to resource, scale and are constantly under the watchful eye of the industry, meaning they're prepared to stay one step ahead of the curve and become 'early adopters.'

Yet on the other, many legacy enterprise data centres were constructed before the current focus on sustainability and, potentially, without adherence to industry standards. So one might argue that, due to the time in which they were designed and built, they have now become inefficient and wasteful, at

least when compared to the latest facilities.

To address this legacy data centre emissions challenge, we believe it's paramount that operators embrace a three-step approach, to achieve the perfect balance of performance, efficiency, and a lower carbon footprint.

The first is to assess your environment – and utilize computational fluid dynamics (CFD) software to help identify inefficiencies within your system. This might include hotspots, poor airflow and/or inefficient rack configurations.

The second is to optimise your existing infrastructure – and partner with best-in-class, global engineering experts to optimize your existing data centre footprint. Here, simple steps may include consolidating mission-critical infrastructure systems, increasing rack densities and improving cooling or airflows to the infrastructure.

The third is to drive data centre performance – through a comprehensive and bespoke modernisation program. One where operators can retrofit or re-design the whitespace with a new containment architecture. This offers optimized performance and greater energy efficiency, alongside reduced operating expenditure (OpEx) and carbon emissions. By enhancing your facility with a simple to follow and data-driven approach, which optimises your existing infrastructure you can achieve an average PUE reduction of 0.4 immediately at the end of the process.

A bespoke approach to modernisation

It's important to note that while these three steps offer a standardised approach to modernisation, the sequence is unique for every data centre.

In fact, there's no one-size-fits-all solution, and once a CFD assessment has been carried out, and the customer is





presented with the resulting environmental impact assessment report, a range of strategies are possible.

The optimise phase is not prescriptive and considers the best way to use a customer's existing data centre infrastructure, wherever possible. Rather, it provides a range of possible improvements that can be made, and the solutions which will help end-users achieve their objectives.

This might mean retaining the current power and cooling infrastructure, for example, but replacing the rack/cabinet layout with standardised units. Further, the cabinet sizes remain, but a containment solution can be designed to optimise the power and cooling requirements.

During the perform stage, customers can really begin to see the benefits from their sustainability and energy efficiency improvements - often via reduced PUE, 20-25% lower energy costs and a decreased carbon footprint. For most customers, the speed and reliability of this end-to-end process, combined with the cost-savings achieved, are the immediate, major

benefits. However, customised solutions, which incorporate cable trays and bus pathways, also provide an increase in available kilowatts per rack and operating temperatures. Suddenly, high performance computing (HPC) is a very real possibility in a legacy data centre.

The good news is that, by improving the performance of your facility with a containment solution, you are also improving its sustainability credentials, and as your PUE reduces, so your energy efficiency increases.

By adopting this three-stage approach, with CFD analysis as the catalyst, you'll almost certainly also have the opportunity to reduce data centre water consumption and be able to address other sustainability issues along the way.

Much of the industry has already committed to challenging Net Zero targets, and to achieve them, we believe legacy data centres need to undergo an intensive energy efficiency improvement process. The approach outlined in this article is a great way to start.

During the perform stage, customers can really begin to see the benefits from their sustainability and energy efficiency improvements - often via reduced PUE, 20-25% lower energy costs and a decreased carbon footprint



François Sterin

Chief Industrial Officer, OVHcloud

ovhcloud.com

Why data is the key to sustainable supply chains

WE ARE PRODUCING and relying on data more than ever before. In 2020 alone, every human created at least 1.7 MB of data per second. As the pandemic motivates even more businesses to migrate to the cloud, the data centre industry has a responsibility to help create a sustainable environment and future for how we manage and store data.

Over the last year, our industry has made significant progress in improving the green credentials of both our own data centres and those of our wider supply chain. The use of sustainable energy sources, liquid cooling and the pursuit of a circular economy model are great building blocks for refining a sustainable business approach in the data centre industry. It is important that we understand the footprint of our physical and digital assets – such as our servers – and what more can be done to reuse components and prolong their lifespan.

For the last 20 years, OVHcloud has been working towards optimising data centre energy consumption. For almost a decade, we have operated a circular economy to create the second and third life of our physical assets. The next steps must focus on building sustainability into supply chains through the expansion of circular economy principles, choice of location and server production. This is all part of turning a company's environmental commitments into tangible and accountable action that positively contributes to mitigating the impact of the climate crisis.

Here are some of the key considerations we are focusing on at OVHcloud, to help achieve our goals of producing 0% waste to landfill and being carbon neutral by 2025, and Net Zero by 2030.

Analyse and optimise consumption

Monitoring energy consumption through PUE (power usage effectiveness) and WUE (water use efficiency) enables businesses to evaluate how these can be managed most sustainably.

At OVHcloud, our water-cooled servers produce 1.1 - 1.3 PUE, which is lower than the industry average. We also stopped using air conditioning in our data centres in 2010, instead opting for natural air cooling with a WUE of 0.24 - 0.29 L / kWh IT (115% lower than the industry average). Continued focus on WEC (water, energy, carbon efficiency) in terms of data centre consumption can also be optimised by achieving ISO certifications 14001 and 50001, which recognise efficiencies in resource use, waste management and energy performance.

Data = power

We have partnered with leading European technology research institute Inria (National Institute for Research in Digital Science and Technology) to enhance our sustainability monitoring and offer the same to our customers. We deploy Inria's real-time energy consumption and carbon footprint data via an API or web portal, enabling users to monitor exactly how they are tracking against their sustainability targets – and identify and address where they are falling short.

Solutions through the supply chain

As well as tapping into the expertise of the likes of Inria, data centres also have the power to positively impact the carbon footprint of their supply chain and wider ecosystem through full life cycle analysis.

This includes developing sustainable packaging and inviting suppliers to sign the Standard of IT Recycling Conduct. Data centres can also help clients address their sustainability challenges by providing real-time energy information at VM level and developing best practices for software energy efficiency. It is the responsibility of all of us within our industry to do as much as we can to embed sustainable practices into our business models and make real commitments to being part of positive change. The future of data centres must be sustainable, and those who refuse to take steps to reduce their environmental impact will simply get left behind.

Leading edge piping systems for data center mission critical cooling

Highly efficient, safe and fast

Benefits

- 25% better energy efficiency
- 100% corrosion free
- Off-site fabrication with bespoke engineer


In today's digital world, data centers are an essential part of our infrastructure, that suffers from an increasing pressure to build quicker, achieve higher levels of energy efficiency, and avoid downtimes during their operations. For that, it's required mission-critical cooling plants where the equipment and piping must be highly reliable and as energy efficient as possible, but also fast in its installation. But, can plastic pipes be used in a data center? What added-value do they bring? Which applica-

tions and where? The Data Center green and brownfield market are under increasing pressure, and no risk should be taken. The Swiss company GF Piping Systems provides leading-edge piping systems for reliable mission-critical cooling. Their range of leading-edge plastic piping systems, which can be used for numerous applications in data centers, cater to these high demands of the industry, offering significant advantages during commissioning and operation.



Mark Stuart
Sales Director Data Centers Europe

Ebnatstrasse 111
8201 Schaffhausen / Switzerland
Phone +44 78 03 242 931
mark.stuart@georgfischer.com



For more information about innovative cooling solutions, calculation tools and other valuable technical data to key engineering questions, sign up in our Data Center dedicated Portal: data-center.gfps.com



Chris Cutler

Business Development Manager, Riello UPS

riello-ups.co.uk

Data centres as partners of the energy transition

HISTORICALLY, our electricity network has operated according to a relatively straightforward premise – to ensure there’s enough supply to meet demand. But with ongoing decarbonisation necessary to tackle climate change and help meet ambitious ‘net zero’ targets, network operator National Grid ESO predicts “the energy system of the future will be based more around supply, with demand adjusting to use or store energy as required”.

While such a fundamental shift towards cleaner and greener power generation inevitably poses plenty of challenges, this transition also offers opportunities too. All data centres deploy uninterruptible power supplies to minimise the risk of damaging service downtime caused by a sudden loss of electricity or fluctuation in the quality of supply. Typically, valve-regulated sealed lead-acid batteries (VRLA) have dominated the UPS market. But more and more manufacturers are now incorporating lithium-ion batteries into their portfolio of data centre solutions. Here at Riello UPS, for example, we’ve introduced dedicated lithium-ion versions of our award-winning modular UPS Multi Power, super-efficient NextEnergy, and transformerless Sentryum series.

Compared to VRLA, lithium-ion batteries have a much better power to weight ratio and far higher power density. This means they can deliver the same power in 50-75% less space. This makes them ideal for edge or distributed IT applications where floor space is at a premium. Or you can simply install extra battery capacity – up to double – in the same footprint. That means you’ve got ample autonomy for emergency backup with the scope to utilise the remainder for more general energy storage and demand response.

Another significant advantage is that li-ion cells can operate safely in ambient temperatures up to 40oC. Contrast this to VRLA batteries, which require regulated temperatures of 20-25oC to perform at their best. This tolerance to higher



temperatures significantly reduces your need for energy-intensive air conditioning and cooling. In some circumstances, it can even eliminate the need for a standalone, temperature-regulated battery room. In addition, lithium-ion batteries offer much faster charge/discharge times and have up to 50 times the number of cycles. To function effectively they do need a sophisticated battery management system that monitors each cell using electronic circuits to maintain balanced states of charge. But this real-time monitoring has the added advantage of enhancing overall system reliability by identifying any early signs of deterioration or weakness.

Using lithium-ion batteries with a UPS has the potential to transform a data centre’s standby power from a purely reactive device waiting to kick in when there’s a blackout into a dynamic energy storage device helping to balance the grid in real-time through mechanisms like peak shaving or frequency response. Not only does this provide tangible economic and environmental advantages for the data centre (i.e. cutting air conditioning requirements; storing cheaper off-peak energy to use during times of peak demand; reduced grid tariff charges; and greater control over TCO). But it would also show the industry is genuinely serious about the wider issue of sustainability and focused on being a key part of the solution, rather than part of the problem.



Tony Craythorne
CEO, Bamboo Systems
bamboosystems.io



Arm-based servers significantly reduce the carbon footprint of data centres

ORGANIZATIONS are increasingly sensitive to the environmental impact of all aspects of their operations, and in particular the fossil fuel required, along with their resulting CO₂ emissions.

Shifting to renewable energy sources is a major initiative being adopted across many industries to address this challenge. However, lost in the discussion is the enormous energy and carbon savings available from adopting more efficient technologies at the outset.

One area with tremendous potential for energy efficiencies are the world's data centres, which today the IEA estimates use about 1 percent of the world's total electricity output.

We recently commissioned an analysis from an independent expert that shows that an Arm-powered data centre reduces CO₂ production by 74 percent, equivalent to that in almost half a million barrels of oil. Let's look at that model, and I will quote liberally from the report which may be found in its entirety here.

Innovative server design can significantly reduce data centre energy use. The implications on CO₂ production from a large data centre shifting to Arm servers is significant. Take, for instance, a model that uses a server's power draw and physical size as the pivotal factors. After calculating server usage, the energy needed for other data centre subsystems can be postulated by using ratios validated by estimates from

Category	annual electricity consumption (MWh)
Servers	251,795
Comms	17,567
Power supply	58,557
Cooling	193,238
Storage	64,412
Total	585,570

data centre operators. The model starts by calculating the energy used by a medium-sized data centre with 750 racks of conventional 1U servers, (42 servers per rack /31,500 total). If we assume that servers comprise two-thirds of the floor space and that the equipment space accounts for 65 percent of the total, this translates to a roughly 62,000 square-foot data centre.

For comparison, hyperscale facilities operated by the major cloud vendors can be a million square feet or more. Knowing server power consumption and the ratios for the various equipment categories, this model uses simple algebra to calculate the energy consumption for each subsystem. From this number, we can derive the usage for each category to be as follows:

Annual Electricity Consumption for 750 rack x86-based Data Centre

Using the same formula, we calculate the amount saved by

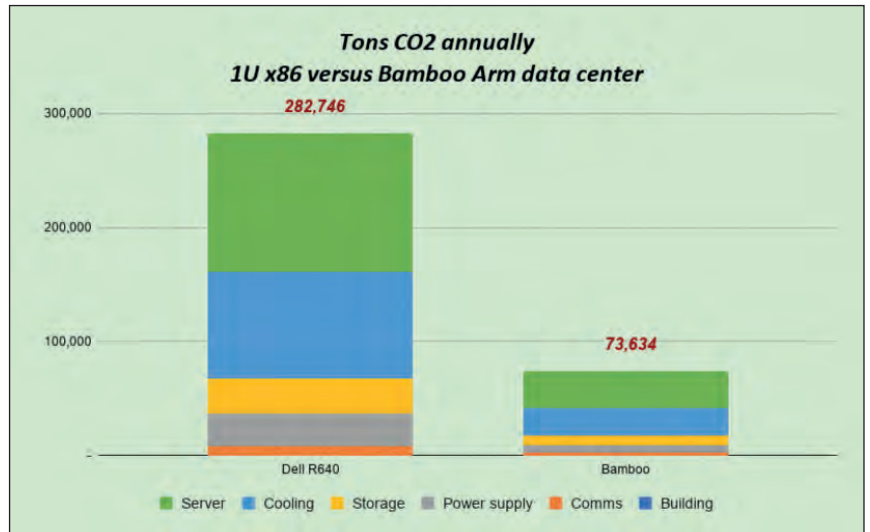
using energy-efficient Arm servers, like our Bamboo Systems' B1000N and compare it to a standard 1U x86 system. Although both are 1U chassis, the B1000N incorporates four Arm compute nodes on a micro-blade, with two blades per chassis for eight nodes per server.

Each server also requires about one-quarter the electricity of a standard server. The algebra shows a Bamboo Arm Server data centre uses only 26 percent as much electricity (152,409 MWh/year) as one based on x86 systems.

The math shows that switching one medium-sized data center to Bamboo can save annually the equivalent to:

- 45,459 fewer cars on the road
- 486,749 fewer barrels of oil used
- 367,076 fewer passengers from JFK to LHR
- 4357 fewer homes with emissions

CO₂ emissions cuts are valuable to the environment, to the bottom line and in the market for carbon offset trading. According to data from the World Bank, the price of carbon ETS



(emissions trading system) credits averages about \$20 per ton CO₂. Thus, transitioning the example data centre to efficient Arm servers could generate more than \$4 million annually. There are many technical and economic reasons to adopt energy-efficient servers in cloud data centres, however, an unsung benefit is their potential for being both good for business and the environment. The next stage for the data centre, is the widespread adoption of the Arm-powered server.






DW ONLINE ROUNDTABLE

BASED around a hot industry topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion.

MODERATED by an editor, this online event would include 3 speakers, with questions prepared and shared in advance.

THIS ONLINE EVENT would be publicised for 4 weeks pre and 4 weeks post through all our mediums and become a valuable educational asset for your company

Contact: Jackie.cannon@angelbc.com










David Watkins

Solutions Director, VIRTUS Data Centres

virtusdatacentres.com



How data centres can help businesses be more sustainable

PROVIDING the infrastructure for society and businesses all over the world 24/7/365, data centres require large energy sources, use a lot of power and generate vast amounts of heat. Data centre power consumption alone amounts to around 416 terawatts, or three per cent of all electricity generated on the planet.

Fortunately, many data centre operators are already realising that their energy consumption cannot continue to rise indefinitely and are committed to making their facilities more environmentally friendly.

Renewable energy

The first step for many providers is a move away from fossil fuels. Data centres are particularly well placed to benefit from renewable energy sources due to their stable power consumption. Indeed, some are already achieving 100 per cent zero-carbon energy in their buildings, resulting in lower emissions of carbon and other types of pollution, as well as cost efficiencies.

A focus on cooling

Cooling is a notoriously energy hungry element of the data centre infrastructure. As much as 40 per cent of power delivered to a facility can be used for cooling and other ancillary functions rather than being delivered to the IT equipment within the data halls. This energy is not being wasted because it is used to maintain conditions of temperature and humidity within, at which the IT hardware can be guaranteed to operate reliably. However, it is not productive energy, and data centre operators are striving to minimise its use.

One way that data centres are becoming more sustainable is the use of outside air cooling instead of high-energy mechanical units. Many servers are now being manufactured that can operate at temperatures as high as 27°C, so natural air cooling

has become a more viable option for data centres in a range of locations.

VIRTUS Data Centres is continually looking at how to optimise cooling. To keep its facilities as efficient as possible, the company uses a variety of innovative design elements for greater efficiency whilst actually lowering costs; this includes air flooded data halls, hot aisle containment, and cooling using a variety of innovative industry leading technologies such as immersive cooling etc.

At VIRTUS' LONDON2 data centre, a borehole was dug at the inception of the site so it can utilise the natural water source for cooling, thereby reducing demand on the mains water supply. Combined with the local climate (if the temperature is below a certain level), the air in the data centre can be chilled without any mechanical cooling, and this efficient cooling technology delivers low Water Usage Efficiency (WUE) for the site.

Return on investment

Green no longer means expensive: tax breaks, incentives and diminishing costs for renewable energy all enable data centres to more affordably and cost-efficiently be green. However, there are more than just cost savings at stake. In the short term, data centre providers that can demonstrate a robust sustainability strategy can enjoy immediate wins, including delivering on government mandates on sustainability, whilst meeting customers' environmental demands.

Increased data can be processed in an energy efficient way, through the deployment of High Performance Computing (HPC) where a smaller footprint of more powerful servers are used, therefore taking up less space and demanding less area to be cooled, using less energy in the process. New and innovative data centre design continues to drive these efficiencies within the sector.



Powering The Connections You Make.

Data centres enable human connectivity. But who keeps the lights on at data centres? We do. Our latest model is ideally suited to data centre demands, offering unrivalled efficiency. **Our products are first rate, but our service makes us really stand out; from pre-sales to in-life support.**

Get a no-nonsense quote backed with deep technical knowledge and a genuine desire to help.

Call 01256 386700,
or visit kohler-ups.co.uk



The new 97.4% efficient 1-6MW PowerWAVE MF1500 DPA

KOHLER POWER
uninterruptible

HYBRID ARCHITECTURE

In-house, colocation and hosting, managed services, cloud, multi-cloud...it's highly unlikely that one-size fits all an end user's requirements. Hence, the hybrid infrastructure models has been gaining traction in recent times.

THE FUTURE IS
HYBRID
WORKING

IN ASSOCIATION WITH



Jan Karremans

Director of Product Marketing, EMEA, EDB
enterprisedb.com

Why cloud native postgres is the natural next step in databases

ACCORDING TO FINDINGS from IDC Global DataSphere, the amount of data created and replicated in 2020 saw unusually high growth due to the increase in people working, learning, and entertaining themselves from home. It was estimated that 1.7 megabytes of new data was created every second for every human on the planet.

Despite the increase in data creation, IDC said that less than 2% of this new data was saved and retained into 2021 – the rest was created or replicated primarily for the purpose of consumption, or temporarily cached and subsequently overwritten with newer data.

Sluggishness and inefficient data management has led organisations to use different types of databases to manage emerging sources of data production. As a result of this anarchy, many organisations lost their single source of truth in terms of accessing data, applying analytics to it, drawing conclusions from it, and making timely business decisions based on data-backed insights. Merging data is both financially draining and time-consuming - making data management one of the largest operational costs within an organisation's IT budget.

This worldwide data explosion is made up of structured and unstructured data, raw data, data-at-rest, data-in-transit and real-time data, all of which navigate IT infrastructure systems at an overwhelming volume, variety, versatility and velocity. This necessitates having the right mechanisms in place to filter and capture the right data sets for the right purposes to extract value, either through analytics, data science or artificial intelligence. But how should organisations embark on this pressing data-driven transformation? Cloud Native Postgres may offer a solution to organisations struggling with data management.

Why is cloud native postgres the next step in databases?

External market forces are urging today's organisations to move faster than ever. Businesses are scrambling to capitalise on new digital revenue opportunities. This involves building new applications, migrating existing ones, and advancing Cloud strategies. Cloud Native Postgres enables this acceleration by offering decision makers the benefits of speed, simplicity, flexibility, and practicality.

Speed: Cloud Native Postgres enables developers to expedite their software delivery projects across testing, staging, and production. Compatibility with Kubernetes makes deploying databases in the cloud less labour-intensive for DevOps teams, allowing IT decision makers to reallocate IT staff to more strategic tasks to go to market faster with the right architecture.

Simplicity: Unlike traditional databases that require external tooling to work in Kubernetes, Cloud Native Postgres works seamlessly in these modern IT infrastructures thanks to its collaboration with Red Hat and Kubernetes. Also, processes such as updates are automated, making the database simpler to manage.

Flexibility: Cloud Native Postgres enables organisations to deploy and manage Postgres databases wherever the business requires them, whether that be in a multi-cloud or a hybrid environment. This type of database easily scales up and down to accommodate fluctuations in business operations, scaling up to accommodate large employee onboarding, customer acquisition, and digital transformation initiatives, while scaling down to respond to internal and external restructuring projects.

Practicality: One of the largest costs for many organisations is the processing, storage, and management of data. With



Cloud Native Postgres, many time-consuming tasks – such as managing external tooling or performing updates – are eliminated. This enables IT teams to win back some of the time spent on data management, and frees them up to focus on critical tasks, making data management more cost-effective.

Cloud Native Postgres offers many benefits to organisations. The technology aligns with organisations' existing needs to modernise and scale increasingly complex IT environments. It also helps IT teams make data management smoother and reduce operational hassles, while helping the business get applications to market faster.

Sluggishness and inefficient data management has led organisations to use different types of databases to manage emerging sources of data production. As a result of this anarchy, many organisations lost their single source of truth in terms of accessing data, applying analytics to it, drawing conclusions from it, and making timely business decisions based on data-backed insights. Merging data is both financially draining and time-consuming - making data management one of the largest operational costs within an organisation's IT budget



Kris Beevers

CEO, NS1

ns1.com

Putting the future of data centres at the edge

EDGE COMPUTING and edge networking have been adopted rapidly in recent years to bring processing near to users and take advantage of lower latency and the opportunity to optimise the routing of data, applications and traffic across distributed footprints.

The drive towards edge computing came from the need to improve application performance and optimise server resources, particularly in the age of IoT and 5G. Among the early adopters were telecom providers who already had a footprint of local 'hubs' serving fixed line connections. Repurposing these has created a network of micro data centres. By partnering with cloud providers, the telcos have brought processing from centralised on-premises or cloud data centres out to the edge.

This, however, has presented various hurdles. Building applications that exist across a highly distributed footprint is quite different from building them to be accessible from just a couple of data centres. To make this approach both scalable and repeatable, companies have turned to the use of containers, which have made edge more accessible. There have been some challenges in how data centres support containers, requiring changes to associated platforms and approaches to networking. Increasingly though, companies are running their workloads across globally distributed Kubernetes clusters – the nodes that run containerised applications – or across serverless functions inside service provider environments.

How one approaches networking will significantly influence the success of edge computing. Classic deployments of micro data centres at the edge aren't useful without leveraging intelligent traffic steering to move data across networks in an optimised fashion. The trailblazers in tech are leading the way with distributed 'edge networks' housed in hundreds of PoP data centres located around the world. DropBox is a great

example. The company has a vast edge network supported by an omniscience infrastructure and has invested a great deal in infrastructure for intelligently orchestrating its application traffic and automating the lifecycles of the resources that support it. They can spin up and control capacity close to audiences served by PoPs in edge data centres to ensure consistent, superior user experiences.

In the world of online gaming, where the minimising of milliseconds and nanoseconds is vital to success, start-up enterprises are pushing their highly performant networks out to the edge, using software-defined networking that intelligently steers traffic on a packet-by-packet basis.

The potential use cases for edge networking are many. Industries such as healthcare can optimise edge networks housed in regional data centres in or near to major cities and serving vast populations. This approach can enable medical diagnostic tools to work rapidly to generate results, improving the diagnosis and treatment of conditions. It can also enable rapid processing for the enormous amount of data generated by healthcare IoT devices, and support network speeds that give medical staff the option to engage remotely with patients without the worry of internet interruptions or delays in the connection.

The development of applications continues to grow, as does the desire to deliver services more quickly, more securely and with more control over data. This is prompting investments in new platforms and foundational services, which allow organisations to build and manage highly distributed edge-enabled networks with data centres that are located close to where both business and consumer demand is greatest. Innovation in these technologies will take different forms, particularly with the emerging dominance of hyperscalers, but will be characterised by the desire to achieve ultra-low latency. The move to the edge will define the future of the data centre from this point forward.

Patrick McFadin

Vice President Developer Relations, DataStax

datastax.com



Distributed computing and Kubernetes: Getting data everywhere

THE FUTURE for building cloud is hybrid and distributed. Everest Group found that 72 percent of respondents want a hybrid-first or private-first cloud strategy. Gartner predicted that 85% of enterprise infrastructure strategies will be based on hybrid cloud by 2025, bringing together on-premises, colocation, cloud, and edge. Bain found two thirds of CIOs already use multiple cloud providers to deliver without lock-in.

Defining what hybrid means in practice will vary, depending on your role and organizational goals. For CIOs, it involves delivering a business-aligned strategy for long-term growth and keeping ahead of competitors. For enterprise architects, translating strategy into infrastructure designs involves specifying the right stack, identifying cloud services, open source data components, and how they work together. For developers, the future is about how to get applications built faster, ready to scale with the ability to create and use data more efficiently.

Making all this work is a challenge. It's not that we don't have the tools already. In fact, we have so many options that we need to concentrate on how to get from ideas to production, and on how to standardise our stacks so that we can make things 'just work.' Rather than having to go into the intricacies of design, we should just be able to build.

Leading the way is Kubernetes, which makes it easier to manage how applications get integrated and deployed. Kubernetes automates all the tricky steps involved in managing capacity, removing friction for developers to get their applications out. It has become the default standard for developers to manage modern application projects.

Originally, it was designed for stateless application components that could be turned on and off as needed, but Kubernetes has been steadily improved to work with stateful workloads like data that may have to exist for years.



What does this mean for hybrid deployments? Kubernetes can support migrating workloads and running them across multiple locations by creating entire virtual data centres. Stateful application components like databases can now be included in the entire service stack. Distributed databases like Apache Cassandra that support multi-region and multi-cloud deployments can enable an active-active data picture across every location. In combination, this gives any organization a lot of flexibility and choice when it comes to where to run your application and avoiding lock-in.

For the developer, this approach lets them build applications that can scale, regardless of the cloud provider they start with and with fewer trade offs. For the enterprise architect, this approach helps them ensure that infrastructure will be able to scale in the right way. And for the CIO, this will all help to deliver that big picture vision around data while remaining flexible to change and new initiatives. The time to plan and do this is now. If you are interested in learning more about how to containerise your data and deploy your entire stack in one control plane, there are some excellent tutorials online. To build this future will take cloud-native apps and cloud-native data too.



Herman Chan
President, Sunbird Software
sunbirddcim.com

5 best practices for edge infrastructure management

ACCORDING TO Gartner, more than 75% of all enterprise-generated data will be created and processed outside of the traditional data center or cloud by 2025. Instead, much of this data will be handled by edge data centers.

As the volume of data from transformative technologies like AI, IoT, and 5G grows, there will be an increased demand for more edge data centers to process and store that data locally. However, managing edge infrastructure poses unique challenges compared to a traditional data center, requiring remote management tools.

ASHRAE advises that “the disparate nature of edge data centers makes the role of remote management and monitoring not only central, but ever more critical, to their operational performance. Tools such as DCIM play a key role in ensuring that all the relevant parameters can be remotely monitored.”

To reduce the complexity of edge infrastructure management, you should follow these proven best practices:

Get the most out of space and power resources

Edge sites are small, so it is critical to maximize the utilization of available capacity. Use DCIM software to visualize and correlate common capacity constraints like cabinet weight, percentage full, budgeted power, and actual power. What-if analysis can be conducted to understand the potential impact of installations and decommissions on capacity, enabling you to identify stranded capacity and defer capital expenditures.

Issue visual work orders to remote hands

Moves, adds, and changes must be completed accurately and quickly, but directing remote technicians on what to do in an edge site can be difficult. Leverage 3D visualizations to see where your cabinets are located and rack elevation views that

provide the exact U-position and other details of your assets. Then, you can show technicians exactly where data and power ports are located so that connections are made correctly.

Monitor data center health of all locations

Leverage an enterprise health dashboard that provides an at-a-glance view of the health and capacity of every single data center and edge site in a single pane of glass. Be the first to know of potential issues that can cause downtime such as power capacity limitations, hot spot formation, loss of redundancy, and other power or environmental alerts.

Accurately track and manage all assets, parts, and connections

Edge sites are difficult to manage, with complex and distributed deployments spanning many sites and multiple business applications. Keep an accurate, real-time inventory of all assets like servers, networking equipment, rack PDUs, patch panels, cabling, and parts and spares. Use visual circuit trace diagrams to avoid overloading circuits, decrease latency, and quickly troubleshoot connections to reduce downtime.

Protect all sites and assets

Securing remote equipment can be difficult without the right tool. Deploy a solution that provides centralized security management including the ability to remotely control electronic cabinet door locks for remote hands, restrict access with role-based permissions, ensure compliance with real-time audit logs, and monitor video surveillance feeds. Edge infrastructure management is a challenge when you have no onsite personnel, no visibility into what’s happening, and potentially harmful environmental conditions near IT equipment. However, by following best practices and using the right software tools, you will improve visibility, uptime, efficiency, and people productivity across all your sites.

Simon Michie

CEO, Pulsant

pulsant.com



Unlocking the edge opportunity: The role of regional data centres

MORE DISTRIBUTED WORKING is placing ever increasing pressure on network capacity as well driving the need for high-speed access to the cloud. Crucially, the evolution of IoT and 5G will accelerate this further, with 5G creating new applications that will place higher demands at the point of use.

As a result of these demands, proximity is going to become key. However, public clouds can be too far from the users and devices that they serve while regional businesses are hindered by being located away from technology hubs. This leaves regional businesses facing bigger latency issues as network traffic has to travel a further distance, which also drives up data transport costs.

All of this is creating demand for decentralised cloud connectivity in those regions that are further away from the major technology hubs. This is then driving the need for edge computing networks which makes the role of the regional data centre even more pivotal.

Regional data centres have a key role to play in developing that architecture and providing the platform for delivering services to the edge, bridging between centralised platforms and micro-edge locations, such as masts and base stations.

The gateway to edge opportunities

While there are clear opportunities that edge computing is unlocking for regional businesses, understanding how to access these may not be so clear cut. This where the need for a nationwide edge computing platforms is evolving.

A nationwide edge computing platform will enable businesses operating across regions to access and take advantage of edge through strategically located regional data centres. Such a platform could deliver critical connections to digital ecosystems, enabling data processing to be carried out in regional and strategic locations and support the latest innovative



applications. Each regional data centre would be configured as a connectivity hub and connected by agile networks that are capable of delivering low latency connectivity across the length and breadth of the nation.

This then places modern compute stacks closer to each digital ecosystem, enabling faster access to these communities for application users, devices and providers. As a result, businesses could have access to distributed multi-cloud models via regional data centres, as opposed to the traditional centralised data central model, breaking down the barriers of location.

Most cloud service platforms will provide at least some distributed cloud services that execute at the point of need by 2025, [according to Gartner](#). However, without a fully operational UK edge computing platform, it's arguable whether organisations will ever be able to truly realise the benefits of distributed edge cloud. A UK-wide platform will be a crucial enabler in supporting the ecosystem of developers which will drive the next wave of innovative applications.



David Walker
Field CTO, EMEA, Yugabyte
yugabyte.com

On Hybrid Architecture

IN, OUT, SHAKE IT ALL ABOUT! One centralized data centre; a central facility, plus regional and or local resources. As edge momentum, fuelled by IoT and 5G, grows substantially, how does data centre architecture need to evolve to ensure data is available to end users where, when and how they need it?

We're now seeing the next evolution in data centres, and it is tied to the move from first generation to second generation cloud deployments, driven by the growth of 5G and IoT.

The first generation was about accepting the use of cloud and getting 'something' into the cloud. As the use of the cloud grows, businesses become locked into the cloud provider; and as IoT drives up data volumes, costs rise.

To break cloud-vendor monopolies, some organisations put specific classes of application in specific clouds, such as Azure for Office, AWS for OLTP and GCP for Analytics. Whilst this addresses some of the cloud provider reliance, it still means being dependent on one cloud provider for each class of application.

The evolution of a second generation of cloud deployments will be characterised by three main attributes.

Firstly, the usage of the same application and particularly databases across cloud providers. It won't be sufficient to run applications in multiple single clouds, they will need to stretch across clouds and on-premises. This will enable organisations to move processing and storage closer to the user and onto the most cost-effective platform for the workload. Organisations using cloud-vendor specific databases, will have to re-think that approach. If you want to migrate your stateless, cloud-native, micro-services across cloud vendors then you need a data layer and in particular a database that is supported cross cloud.

The second trend will be the use of three data centres, across regions, rather than two. The use of two DCs, one active, one passive is replaced by distributed databases in three locations.



If you lose one, the services will still survive. Running in Frankfurt, London & Dublin is no more difficult than running in any two of those. It's simple and more resilient; and it's easily adopted with the technologies now becoming available. Three data centres will also help with data compliance and security. For example, having data reside exclusively inside only one of the US, the EU and the rest of the world to meet regulatory requirements whilst also being in a single global database is now possible.

The third trend is the move to genuinely agnostic APIs for data storage. I predict that the database will settle on the Postgres API for SQL databases and Cassandra for No-SQL, whilst block storage will look increasingly like S3. This de-facto unification of data storage APIs is important because it removes a barrier that can make scaling difficult and prevent organisations moving between data-layer platforms.

Commoditisation of the data centre is already happening. These three trends are the inevitable outcome of the desire to cost-effectively scale without being tied to specific vendors. The vendors won't give up their lock-in advantages without a fight, but it will happen. This evolution will bring plenty of innovation, not only at the edge, but also critically in how we build our data centres of the future.

COLOCATION + HOSTING

In-house or to outsource? Colocation and hosting facilities offer a range of significant advantages when it comes to accessing data centre infrastructure. For many end users, these advantages are far too compelling to ignore: the very latest infrastructure, constantly updated, constantly supervised, with constant access to expert advice and support for a fixed monthly fee.



IN ASSOCIATION WITH



Matt Parker
CEO, Babble
babble.cloud

Adopting a hybrid model

THE FASTEST GROWING, most efficient and resilient businesses usually have one thing in common – they have all implemented cloud-based technologies. From start-ups to large-scale corporations, the benefits of cloud-based communications are undeniable and offer organisations a competitive edge against others.

A new normal - adopting a hybrid model

The past 12 months have made it abundantly clear that we're on the cusp of a whole new way of working. The future of office life for many organisations will be a hybrid model, balancing working from home with office-based activities that can be completed more effectively through face-to-face interaction, for example for meetings and brainstorming. For those adopting a hybrid model, effective use of cloud technologies will be an

essential tool in the armoury - facilitating agile working and collaboration amongst employees, as well as opening up endless opportunities for a company's workforce.

The ability to get online to access files and emails at any time or place has never been more important, and those that don't adapt will pay the price when it comes to productivity and business growth.

Cloud reduces costs

Cloud-based solutions reduce operating costs for businesses. Replacing the need to purchase expensive hardware with simple and transparent pay-as-you-go subscription services means that companies can reinvest their profits more strategically. The return on investment alone should make



business leaders seriously consider switching to a cloud-based solution. If a business is hoping to rapidly grow and evolve over the coming years, then they'll need to keep abreast of technological trends and advancements in order to thrive.

Prioritising customer service

Keeping customers happy in 2021 has never been more important, and those that maintain high standards of customer service will be ahead of the game. Transforming from a legacy contact centre to having all inbound and outbound customer communications handled in one central cloud-based solution allows businesses to talk to more customers and have better, more meaningful, and ultimately profitable conversations.

This is something we've seen first-hand with Yopa, a full-service estate agent, that we helped transform. The automated contact centre solution that we implemented syncs all call data, ensuring agents are connected with specific leads that are most likely to generate sales, before moving on to dial the next lead with the best conversion potential. As a direct result, the company has seen a 75% increase in monthly talk time, improving from 400 to 700 hours of productive talk time over the course of the year.

Secure, flexible solutions that help your business thrive

This year, businesses must minimise operational challenges

through implementing solutions that offer complete flexibility. It doesn't matter whether you're a start-up or large enterprise, cloud solutions can be scaled up or down depending on the ever-changing demands of a business. Not only does this give business leaders a competitive edge, but it also minimises the risks related to operational maintenance issues while saving on the costs of changing and purchasing physical infrastructure.

If 2020 taught us anything, it's that businesses need to be ready for the unexpected, and cloud-based solutions have shown that they're the most effective way of achieving this. It doesn't matter how big your company is – security breaches can affect anyone. Authentication access control or data encryption are just some examples of how data can be protected in the cloud, but any businesses leader looking to cement a competitive edge mustn't let cybercrime set them back on their path to success.

The need for businesses to be more resilient is more apparent now than ever before. Whilst the future's unknown, leaders need to ensure that they don't have their heads in the sand when it comes to reviewing their business operations and technology. In order to achieve long-term strategic goals, offer protection and increase efficiency, businesses must make the switch to cloud-based solutions before they're left behind.



DW **DIGITALISATION WORLD**
Modern enterprise IT - from the edge to the core to the cloud

New product and process development is the foundation for the growth of the DW industry.

If you want to highlight the recent important breakthroughs that your company has made, please submit an abstract to philip.alsop@angelbc.com

It is imperative that Digitalisation World Magazine remains a timely resource for this industry, so we are especially interested in highlighting very recent work.

IN ASSOCIATION WITH



Ben Parker
CEO, IDS (An STL Company)
impactds.com

Data centres are in transition, what choice will organisations make?

EVEN BEFORE THE PANDEMIC, changes in the enterprise data centre market were afoot. Eighteen months of unprecedented demand for data have only served to accelerate the expansion of the hyperscalers. This is forcing enterprises to make some serious choices about their network infrastructure strategies for the future. Do they opt to move to the cloud, stay on-premises, or move into a colocation data centre?

With the working world moving towards a hybrid model, the large city-centre office footprints occupied by enterprises may no longer be necessary or make economic sense. When enterprises divested their on-premises data centres in the past, the natural choice was to move to a colocation facility with all the benefits of interconnections to SaaS and PaaS solutions allowing them to scale with little complexity and controlled costs. The options now, however, are broader as the lines delineating the data centre space have become blurred.



The massive demand for capacity has prompted the hyperscale players to spend billions building and expanding their global data centre footprints. One recent example is [Google](#), which has bought land just north of London on which it can potentially house its first owned facility. Meanwhile, just last month it was [revealed](#) that AWS has bought the site of the disused power station in Didcot, Oxfordshire with the intention of building two hyperscale data centres.

Facilities on this enormous scale, however, take time to create and bring colocation data centres into play. According to [research](#), 70 per cent of the world's hyperscale data centres are located in facilities leased from data centre operators such as Equinix. For colocation developers, it is easier to build facilities and hand over the keys to one organisation than it is to manage leases for multiple smaller enterprises. The rewards are so great that they are prepared to take the risk of building the facility even before they have decided whether to manage it or sell it. What this means for the enterprise is that they now have far more choices and no longer have to own a data centre. If workers return to the physical office over the next few months, more space will be needed to ensure social distancing can be implemented. It makes sense, therefore, to sacrifice the traditional on-premises facility for the efficiency, security and scalability offered by the cloud and/or a colocation data centre, regardless of who owns it.

This market, however, is fast moving, and the conversation is already moving beyond the evolution of the colocation facility to edge. As digitalisation expands further and smart technology inside and outside the home is widely adopted, the need for edge computing power will grow. This will create a demand for edge data centres that can deliver processing power close to where it is most needed. The expectation is that the key colocation and hyperscale providers will corner this market sooner rather than later, adding further choice and complexity to the landscape.

Jack Bedell-Pearce

Co-founder & CEO , 4D Data Centres

4d-dc.com



Colocation in a COVID world

THE DEMAND FOR COLOCATION capacity has been rising steadily over the course of the past few years, as enterprises look to decrease their reliance on private, in-house data centres to host their information. Businesses have also become increasingly digital and as a result more reliant on data and tech to be successful.

The current pandemic has not slowed demand. If we think of colocation centres as the epicentre of connectivity, data storage and processing, as well as a variety of business-critical applications, it shouldn't come as a surprise that these hubs have thrived in a year when we've all been homebound and more heavily reliant on devices to keep connected and productive. It is also only natural to believe that data storage and internet use will continue to increase in the years ahead. Lately there has been a stream of colocation and connectivity projects at an unprecedented rate. The sector is expected to register a CAGR of 5.5% between 2021 and 2026. COVID-19 has in many ways facilitated the acceleration of colocation trends, prompting colocation providers to think differently about how they operate and staff their sites to ensure they are using their resources as efficiently as possible.

Colocation becomes a good option as infrastructures expand. Paying for the optimum power and cooling for an on-premise server room can get expensive. Conversely, data centres benefit from setups that make colocation by far the more affordable option, this certainly shows no sign of abating as we move away from the worst effects of the pandemic.

A sustainable alternative

For the data centre industry, which is single-handedly responsible for at least 1% of global energy consumption, sustainability takes on a more intense and innovative path. Based on the sheer size and scope of its business, data centres, like enterprises, have an obligation to implement and promote more sustainable choices and solutions. Colocation providers offer greater scope to introduce new technologies to improve the efficiency and sustainability of operation. For example, immersion cooling - a process which sees computer



components or full servers immersed in a dielectric liquid that enables higher heat transfer performance than air - represents a much more practical measure for addressing challenges around inefficient energy sources. With an ever greater emphasis on methods such as immersion cooling, we will see a larger emphasis placed on this in the next 12 months. It is vital for businesses in the colocation market to plan for eco-cooling technologies so they don't fall behind.

On a cloud

The growth of cloud-based software and storage will in many ways help boost colocation's own continued development. As more people make use of public cloud services, the more data centre capacity the hyperscalers need to rapidly acquire to accommodate demand. The London and South East colocation market is one of the biggest beneficiaries in this sphere. London is home to the largest data centre market in Europe and the UK has been ranked as one of the top places to build data centres, indicating the importance of the UK for the global data centre industry. Many platform as a service (PaaS) and infrastructure as a service (IaaS) cloud service providers are deciding to pursue the colocation route, opening up new sites with colocation providers closer to hyperscale data centres, rather than building and kitting out their own sites, as they expand operations globally.



Herman Chan
President, Sunbird Software
sunbirddcim.com

8 ways to remotely manage your colo infrastructure

ACCORDING to 451 Research, almost 50% of enterprises plan on increasing their use of colocation services over the next two years.

With the use of colo data centers on the rise, data center managers must find ways to remotely manage their assets in colo spaces to maintain uptime, increase the efficiency of capacity utilization, and improve the productivity of people. Some of the most beneficial ways to improve remote colo infrastructure management include:

Monitor, report, and trend your power loads

Monitoring power is critical and knowing the exact load for every rack is important to maintaining uptime.

Leverage an enterprise polling engine that can collect all the power data from inlets, outlets, and circuit breakers, retain that data for long periods of time, and transform it into actionable information. Set thresholds and alerts to be the first to know of potential issues before they become real problems.



Monitoring power is critical and knowing the exact load for every rack is important to maintaining uptime. Leverage an enterprise polling engine that can collect all the power data from inlets, outlets, and circuit breakers, retain that data for long periods of time, and transform it into actionable information.

Monitor the environment in real-time

Deploying and monitoring environmental sensors allows you to identify hot spots and overcooling, ensure a safe operating environment for IT equipment, reduce energy costs, and improve uptime.

Accurately manage all assets, parts, and spares

Enable a single source of truth by maintaining an accurate inventory of the data center assets across your entire colo deployment including servers, storage, networking equipment, rack PDUs, and patch panels. For more granular asset management, be sure to track all parts and spares like hard drives, memory modules, power supplies, patch cables, and even boxes of screws.

Simplify space and power capacity planning

Leverage business intelligence dashboard charts and visual analytics that your management can understand to show exactly what your current capacity is and when you can expect to run out. Consider a modern Data Center Infrastructure Management (DCIM) solution that can automatically calculate power budget profiles for each device instance based on how they are used in your environment, allowing you to get up to 40% more out of your existing cabinet resources.

Visually document power and data connectivity

Good network documentation practices improve uptime and the speed at which you can deploy equipment. Plus, if there is an issue you can quickly trace cables via your documentation to determine the root cause rather than having to manually trace cables. Poor documentation often leads to mismanagement of resources, ineffective use of capacity, inability to move or install equipment, unsafe operating environments, and increased costs for new cabling and hardware.

Track and manage port capacity

As colo environments grow increasingly complex and dense, it becomes crucial to accurately inventory and track individual

physical port types residing on every device. This helps you make the most informed decisions when planning for new equipment or when maintaining and troubleshooting colo infrastructure.


Measure cable lengths before deploying

Keep your data center free of loose cables, reduce the cost of wasted cable, and save time by not having to visit the data center with a tape measure. With modern data center management software, you can accurately measure cable lengths virtually on your floor map.

Enhance data center security

Protect your colo equipment from unauthorized access to mitigate against physical security threats. Use electronic door locks and a centralized data center security management solution to ensure only authorized personnel and remote hands are opening cabinets. In addition, a real-time audit log shows you exactly who did what and when.

Don't let colo infrastructure management be a challenge. By deploying the right remote management tools and following proven best practices, you can dramatically simplify the management of all your colo deployments.





Terry Storrar
Managing Director UK, Leaseweb
leaseweb.com

Is colocation a feasible option for your business?

COLOCATION SERVICES have moved on from the basic warehouse-type spaces of a few years ago; they are a hot area to watch. Now a viable option for the smallest businesses through to corporates, colocation services are set to be more sophisticated and diverse as they mature over the next few years.

According to a report from ResearchAndMarkets released in May 2020, the colocation market is set to almost double by 2025 to reach a projected value of \$58.3 billion USD. This report also predicts Europe will be a growth hotspot in the next five years, to satisfy increasing demand from hyper-scale companies in the region. The global growth of colocation services is driven by the need for businesses of every size to manage spiralling volumes of data and is also fuelled by factors such as the pandemic, 5G and AI computing-as-a-service. These are changing workforce operations, systems management and business requirements for secure, low latency, high-powered technology across all types of business. Colocation is no longer a consideration for large organisations alone.

The benefits of colocation

Using a colocation service lowers the Total Cost of Ownership (TCO) for the business. It reduces the cost of utility bills, infrastructure and physical space, freeing more time for staff to concentrate on company specific goals and innovation.

Using a modern colocation service can also be a simple way of reaping the benefits of the latest technology without upfront costs, securing a competitive edge while avoiding capital expenditure in favour of a manageable monthly expense. Advanced colocation services also offer distinct benefits including robust security, enhanced business continuity, scalability and expert support. All this can make an attractive alternative to on-premise hosting. At the least, there are compelling reasons for a business to use a combination of on-site and colocation services.



Some industry trends to watch

AI computing-as-a-service will be a key industry driver moving forwards, with low latency and 24/7 support needed for supporting AI assistants and software updates. Data centres have improved and increased their remote management tools, so they are poised to take on an increasing volume of customer operations functions in the next few years. As we emerge from the pandemic, many businesses may downsize their office space and colocation offers a flexible option for a remodelled IT infrastructure to support off-site working. Super-fast 5G connections to data centres will also speed up connections for workers wherever they are located.

Although a snapshot view, this does indicate the pace at which colocation services are adapting to fully support today's demanding business data environments. Gartner's research paper entitled "Your data center may not be dead, but it's morphing", predicts that by 2025, 85% of infrastructure strategies will integrate on-premises, colocation, cloud and edge delivery options, compared with 20% in 2020. No longer limited to the realm of corporates, the colocation services of today and tomorrow provide cost effective choices for any business and are worth considering in current business plans.

NO GREY AREAS

When it comes to critical infrastructure there is no room for uncertainty



REVIEW



ADVISE



PLAN



DEPLOY



MAINTAIN



With over 26 years of experience in the supply and delivery of backup power, we have developed a proven approach to deploy solutions from single units to multi megawatt systems that achieve all technical and commercial drivers.
info@powercontrol.co.uk | 01246 431431



Darren Watkins

Managing Director, VIRTUS Data Centres

virtusdatacentres.com

Colocation is forever

THOSE OF US WORKING in technology have long understood that data centres form part of the backbone to the digital economy. However, the Coronavirus crisis saw the industry thrust into the public eye on a much wider scale. Over the course of 2020, data centres were recognised for their crucial service provision in maintaining business continuity and communications in the face of an unforeseen and particularly challenging period. Organisations have realised that they can't take any risks with their data centre strategies and that demand for space and power will continue to grow.

Powering budget and growth

Onsite servers can be one of the biggest money drains for a business. From employing expert data centre managers and staff, to powering servers ensuring they're kept at a temperature, onsite data centres demand plenty of budget and ongoing Operational management. Also, necessary power and associated cooling for the IT hardware to support such applications just isn't reliably available in many of the office buildings in which businesses are working.



Organisations have found that by moving onsite IT infrastructure to a colocation provider, they can reduce costs, scale up or down, benefit from more uptime and mitigate security and compliance risks – all without any detrimental effect on performance.

Ensuring reliability, uptime and optimum data security Reliability has always been one of the biggest arguments in the build vs. buy debate. Businesses rely on critical applications and core software systems in order to operate at peak performance. But it's not as simple as it appears. Power outages do happen, servers can fail, and technicians can fall victim to human error. However, by collocating to a shared data

centre, companies benefit from a professionally managed environment and an expert team of Operational Excellence whose job it is to keep facilities up and running.

Colocation in a cloud-based industry

When it comes to threats to colocation, there is no question that the rapid growth of cloud has shaken up the industry. But, despite existing misconceptions, the cloud and colocation offerings available today exist in a topology that can be modelled to suit any corporate business objective, including rapid growth, consolidation, simplification, mobilisation or cost reduction.

To remain competitive, traditional colocation providers must ensure that their data centres are capable of delivering not only first-class colocation services with flexible contract options, but also offer their customers a clear on-ramp to cloud. Businesses use public clouds for access to huge amounts of data and massive compute capability, for on-demand computing when needed, or simply for storage.

However, organisations can still maintain their own private clouds as a way of processing and adding value to their own sensitive data, which they collect to handle complex computations. This is the hybrid world that is becoming the de facto standard, and this is how the data centre industry is evolving to support their customers.

It appears that the much talked about demise of colocation was greatly exaggerated. While some anticipated colocation would lose out to cheaper cloud computing solutions, analysts are now universally predicting its continued growth throughout this decade and beyond. The key is ensuring that your provider can meet the needs of today, scale for tomorrow (up and down) and be ready for any unforeseen circumstances in the future - especially now we have experienced a pandemic. By moving onsite IT infrastructure to a colocation provider, organisations can reduce their costs, benefit from a more stable uptime environment and mitigate security and compliance risks – all without any detrimental effect on performance.

AI / AUTOMATION

Artificial intelligence (AI) seems to be the buzz technology right now. AI in all its guises is placing increasing strain on both data centre capacity and infrastructure, as more and more applications and workloads require more, and faster, compute, network and storage resources. And, of course, AI is also finding a place within the data centre, in terms of providing a degree of intelligence and automation which helps to ensure facility optimisation.



IN ASSOCIATION WITH



Ed Ansett

Chairman and Founder, i3 Solutions

<https://i3.solutions>

Dynamic workload provision demands automated and adaptable data centre power provision

EVEN BEFORE the multi-cloud era the holy grail for big IT was the ability to move workloads across different infrastructure platforms. Today, all companies, from hyperscale cloud providers, big SAAS and IAAS providers to enterprise players strive to make applications, dev and ops more dynamic – i.e., host workloads in the best place and move it as necessary.

The goal is to literally abstract the workload from the infrastructure. The open source orchestration revolution of

Kubernetes has propelled much of the progress in this arena, e.g., by enabling the world of flexible Cloud SLAs. However, it is not the only catalyst and different approaches are being tried with firms throwing vast amounts of resources and R&D at the problem. For example, some hyperscalers say they can dynamically move workloads to locations where renewable energy is available to power their data centres. But creating fully dynamic environments where workloads move automatically is a highly complex undertaking and no-one has so far cracked this



IN ASSOCIATION WITH

As a general rule of thumb, data centres don't tend to be operated utilising above 70 percent of their total power. In addition to this headroom maintained for reasons of, e.g., redundancy, some data centre power can also become stranded, or trapped as a result of IT deployments when equipment is added, shifted or replaced in the racks

conundrum at scale. The objective is clear, but even as strides are made, how data centre power can match this dynamism needs consideration.

In almost all data centres, power is literally the life blood of the operation. However, the power topology is fixed at the design stage and is uniformly delivered throughout the facility. In an age of dynamic environments, the question arises whether power can be made more adaptable in order that data centre power chains can be responsive to accommodate multiple IT SLAs?

As a general rule of thumb, data centres don't tend to be operated utilising above 70 percent of their total power. In addition to this headroom maintained for reasons of, e.g., redundancy, some data centre power can also become stranded, or trapped as a result of IT deployments when equipment is added, shifted or replaced in the racks. This gives rise to a second set of questions; is possible to capture and divert this inherent data centre power to where it is needed? Can this be done without adversely affecting redundancy? And is it possible to ensure that flexibility does not come at a price of increased risk of disaster – i.e., to ensure disaster avoidance when the unthinkable happens and grid power is lost?

In short, is it possible to use automation as a means to direct power to the most critical applications and workloads when outage disaster strikes? Can a data centre power system be designed incorporating automated dynamic provisioning where power for environments is regulated to meet SLAs using adaptable redundancy?

The answer is yes. The answer is ARP.

Adaptable Redundant Power (ARP) topologies can deliver exactly this. And more. It is based on four operating modes, Adaptable Redundancy; Inherent Redundancy; Adaptable Inherent Redundancy; Disaster Avoidance.

Adaptable redundancy means companies can achieve the goal of automatically powering applications that are dynamically provisioned as containers even where applications have different SLAs in the same data centre. It does this by overcoming the power system single SLA which inevitably means the power system is either under provisioned or over provisioned. Adaptable redundancy is the ability to dynamically modulate power system availability to align with variable IT

SLAs. It provides variable redundancies, 2N, N, N+1 and N to different cabinets, rows or data halls as needed. Inherent Redundancy is the ability to access stranded power that already exists in the power chain and to use it to provide virtual redundancy. For example, in a data centre comprised of six 1MW halls, the power drawn will typically average around 60% of the power available and in almost every case never peak above 80%. Inherent redundancy enables that stranded accumulated power capacity to be directed as needed with no physical modifications required to the power system.



When everything is operating as normal in a data centre all applications are equal. That is until there is an outage. ARP disaster avoidance means power can be automatically directed to the most critical workloads in the event of extended power loss. In any outage a partial power system failure resulting in reduced capacity can cause a total data centre outage. Using ARP, disaster is avoided as it prevents a cascade data centre failure by preserving power to the most critical applications.

Be adaptable and automated

The pressures on data centre operators to deliver increased sustainability by making better use of existing resources, to be more responsive and flexible with power provision to match workloads moving around while maintaining current levels of resilience and reliability are real. Adaptable power provision is already widely used in other sectors such as critical processes and energy. The key is automation and it is time for the data centre sector to be pragmatic and look at the benefits of Adaptable Redundant Power.



Herman Chan
President, Sunbird Software
sunbirddcim.com

Automation via integration: 3 real-world examples

DATA CENTER MANAGERS are increasingly looking for ways to automate tasks to save time and improve data accuracy.

We've spoken with countless data center experts on how they leverage Data Center Infrastructure Management (DCIM) software and APIs to drive automation. Recently, some of the industry's best from eBay, MacStadium, and the University of Chicago shared their use cases on how they simplify data center management with automation.

eBay automates power budget calculations and updates

Before deploying DCIM software, eBay had difficulty identifying their power utilization for capacity planning purposes, especially for new devices and applications.

Then, they leveraged their DCIM tool's Auto Power Budget feature, and everything changed. Auto Power Budget is a machine learning algorithm that automatically calculates and updates power budget profiles for each device instance based on how they "actually" consume power in their environment. With Auto Power Budget, eBay can easily quantify their return-on-investment. For example, for a project they previously thought required six cabinets in six locations, they now know with Auto Power Budget that it only requires four cabinets in those locations, a 33% savings.

"A server ready cabinet... is about \$10,000 to us," said Ken Torres, Global Data Center Engineer. "So just for this very project... that's a cost avoidance of about \$120,000. It's pretty significant. It's one of those nice things you like to raise up to your directors or senior managers."

MacStadium Automates back-office processing

MacStadium is always working to improve their online provisioning experience. They needed a solution to help manage the high volume of changes they experience without disrupting existing customer experiences or internal workflows. They also wanted to eliminate manual data entry and the possibility of human error.

By leveraging their DCIM software's APIs, MacStadium can now assign slots and assets to be ready on demand for customers signing up on their website. The integration allows them to automatically move assets from their pre-staged account directly to their customers' accounts and updates all their internal systems.

"Using the API, we're able to poll our existing racks, see where we have space available, and assign that space automatically to a customer order via our website," said Robert Perkins, Lead Infrastructure Engineer/Architect. "We sync the rack slots, switch



port, PDU port, make and model, serial number, and asset ID between all of our systems. Changes across our billing platform, customer portal, administration system, and accounting systems are all pulling data from [our DCIM software].”

The University of Chicago is making DCIM the “Center of the Universe”



The University of Chicago has many different tools and systems, and they need to have them all communicate so they can fully understand what is going on in their environment.

“What we’re trying to do as an organization is to get out of the data world and into the information world,” said Raymond Parpart, Director Data Center Strategy & Operations. To achieve this, they are building an entire automation architecture with DCIM software at the “center of the universe for the data center.”

By integrating their DCIM with Slack and Teams, they have created Command Center channels on both platforms, and anybody in the organization who wants to know what’s going on can simply check those channels. Next in their plan is to have those alarms sent automatically to ServiceNow to create an incident ticket.

“We are the colo for the university,” Parpart said. “Anybody here at the university who’s got a system can put their systems in. We’ve got to communicate with all our researchers on what’s going on in their environments.”





DW ONLINE ROUNDTABLE


BASED around a hot industry topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion.


MODERATED by an editor, this online event would include 3 speakers, with questions prepared and shared in advance.

THIS ONLINE EVENT would be publicised for 4 weeks pre and 4 weeks post through all our mediums and become a valuable educational asset for your company

Contact: Jackie.cannon@angelbc.com







IN ASSOCIATION WITH



Peter Ruffley

CEO, Zizo

zizo.co.uk

Utilise AI or risk losing out to the competition

OVER THE COMING YEARS, we are going to see a tremendous investment in large scale and High-Performance Computing (HPC) being installed within organisations to support data analytics and AI. At the same time, there will be an onus on data centre providers to be able to provide these systems without necessarily understanding the infrastructure that's required to deliver them or the software or business output needed to get value from them.

There's no denying that the majority of data centres are now being asked how they provide AI solutions and how they can assist organisations on their AI journey. Whilst organisations might assume that data centres will have everything to do with AI tied up. Is this really the case? Yes, there is a realisation of the benefits of AI, but actually how it is best implemented, and by who, to get the right results, hasn't been fully decided.

Solutions to how to improve the performance of large-scale application systems are being created, whether that's by getting better processes, better hardware or whether it's reducing the cost to run them through improved cooling or heat exchange systems. But data centre providers have to be able to combine these infrastructure elements with a deeper understanding of business processes.

It's easy to go down the route of promoting that 'we can save you X, Y, Z' but it means more to be able to say 'what we can achieve with AI is..X, Y, Z'. Data centre providers need to move away from trying to win customers over based solely on monetary terms.

When it comes to AI, there has to be an understanding of what the whole strategic vision is and looking at where value can be delivered and how a return on investment (ROI) is achieved. What needs to happen is for data centre providers to work towards educating customers on what can be done to get quick wins.

There are some fascinating innovations already happening, where lessons can be learnt. In Scandinavia for example, there



are those who are building carbon neutral data centres, which are completely air cooled, with the use of sustainable power cooling through solar. The cooling also comes through the building by basically opening the windows. There are also water cool data centres out there under the ocean.

As the global costs of energy rise, and the numbers of HPC clusters powering AI to drive our next generation technologies increase, new technologies have to be found that lower the cost of running the data centre, beyond standard air cooling. It's great to see people thinking outside of the box on this with, with submerged HPC systems and full, naturally aerated data centres, but more will have to be done (and fast) to meet up with global data growth.

The appetite for AI is undoubtedly there but for it to be able to be deployed at scale and for enterprises to see real value, ROI and new business opportunities from it, data centres need to move the conversation on, work together and individually utilise AI in the best way possible or risk losing out to the competition.

Oliver Goodman

Head of Engineering, Telehouse
telehouse.net



How AI and automation will impact data centres in 2022

AI is beginning to revolutionise the data centre industry. As it continues to advance, more data centre operators will see its tangible benefits.

Over the next 5 years, we can expect automation to play a more crucial role in improving operational efficiency, sustainability and cyber resilience. As there is a growing pressure on data centres to become greener, AI will become an enabler of change, helping operators drive practices which benefit the environment whilst gaining a competitive edge.

More efficient workload distribution

There is a significant amount of IT workload management that takes place among larger data centre customers to make sure they are using server capacity as efficiently as possible. Data centre selection is still broadly done on proximity to carrier connectivity rather than cost of facility. However, over the next year, AI will play a bigger role in automating the movement of IT workloads in real-time, not only inside a traditional data centre but also in the increasingly talked about edge markets and in a hybrid-cloud setting. IT workload relocation will help bring overall electricity costs and carbon footprint down and free up operators metro sites for the more critical low latency services.

Driving sustainability

As data traffic increases every year, so does energy consumption. AI (and with it machine learning) has the biggest potential to improve energy efficiency in a data centre environment where a significant portion of the energy is utilised on cooling systems. Most data centre control systems already use AI to an extent to turn off redundant modules and put them into hibernation or standby where appropriate, ensuring optimum efficiency for the actual load at any given time. We can expect more data centres to implement intelligent control and cooling solutions in the next few months as the systems themselves get smarter.



AI will also start to play a key role in managing the lifecycle of equipment and parts. With predictive failure and maintenance, data centre operators will know exactly when and which part is likely to wear, safeguarding the efficiency of each component and minimising capital expenditure.

There is, however, a growing need to develop highly efficient systems that leverage automation and intelligent controls to help data centres achieve maximum efficiency gains. Whilst that will not be a quick fix, especially in legacy sites, ideas and strategies for how to achieve optimised efficiencies are well underway throughout the data centre industry.

Better protection against cyber attacks

Data breaches and cyber attacks are now making the headlines on a seemingly daily basis. Four in ten UK businesses reported having breaches and attacks in the last twelve months.

With hackers finding more sophisticated methods to breach networks, data centre operators face the ongoing challenge of implementing stronger cyber security practices.



So far, the developments in network management, AI and cyber security have allowed more data centres to detect a greater number of malicious activities. The advancements in AI will continue to accelerate, meaning more robust solutions that can detect threats are now appearing on the market.

These solutions identify deviations from the typical traffic behaviour and automatically shut down ports on the networks where anomalies occur.

As AI continues to learn the usual patterns and flag suspicious data inflow and outflow, the focus will be placed on speeding up the detection and response processes on the ever-increasing attack surface even more. In addition, better AI integration with network monitoring is on the horizon, for example when it comes to detecting changes in server behaviour after a person enters a secure server hall.

From a cyber security perspective, AI offers virtually limitless possibilities and it will add an additional layer of robustness to data centres. Alongside better security, there is also enormous potential for improving workload distribution and driving sustainability with automation and more intelligent solutions. Utilising AI into the future will help operators keep data centres running more efficiently, facilitating further growth.

So far, the developments in network management, AI and cyber security have allowed more data centres to detect a greater number of malicious activities. The advancements in AI will continue to accelerate, meaning more robust solutions that can detect threats are now appearing on the market. These solutions identify deviations from the typical traffic behaviour and automatically shut down ports on the networks where anomalies occur.



DCS ONLINE ROUNDTABLE



BASED around a hot industry topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion.

MODERATED by an editor, this online event would include 3 speakers, with questions prepared and shared in advance.


THIS ONLINE EVENT would be publicised for 4 weeks pre and 4 weeks post through all our mediums and become a valuable educational asset for your company

Contact: Jackie.cannon@angelbc.com



DESIGN + INFRASTRUCTURE

The design of a data centre, along with the infrastructure - in particular, the power and cooling solutions - deployed, has a major impact on both the environmental and financial performance of the facility. As the IT workload requirements respond to digital demands, so the data centre infrastructure needs to mirror these changes.



IN ASSOCIATION WITH



Stuart Crump
Director of Sales, Iceotope
iceotope.com

If you can't stand the heat...

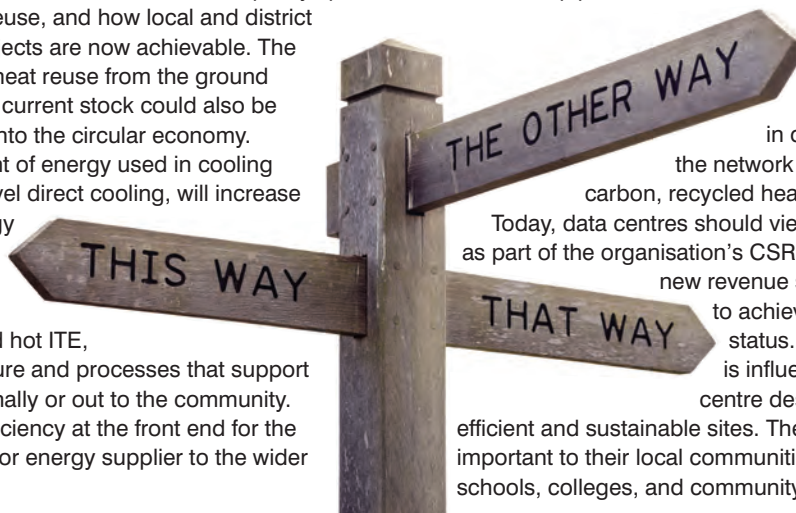
DATA CENTRE and infrastructure design is at a crossroads, being driven by business requirements on the one hand and legislative and societal demands on the other. When did the industry's purpose change from providing resilient data availability to its client base at a reasonable margin, to carrying the mantle of environmental, social and governance (ESG) programmes which demonstrate good global citizenship?

This sea change means our industry is now scrutinising old practices and becoming more receptive to new process designs and methodologies. A key area of design examination is in cooling, and rightly so, as currently the vast majority of data centres are air-cooled, which for the most part requires compressor based chilled systems to force air across the hot IT equipment (ITE). These systems are expensive to procure, have large physical and environmental footprints, require around 35-percent of the data centre's total power budget to operate and are inefficient at cooling ITE. A combination of these factors is why the average data centre has a 1.67 PUE, not only that, but they offer limited opportunity for heat recycling.

I highlight this situation so as to discuss the often quietly spoken about capability of heat reuse, and how local and district heating infrastructure projects are now achievable. The industry must design for heat reuse from the ground up, although much of the current stock could also be converted to bring them into the circular economy. Harnessing the 35-percent of energy used in cooling and investing in server-level direct cooling, will increase the ITE share of the energy budget providing more compute capacity to customers. Direct cooling of the servers and hot ITE, introduces the infrastructure and processes that support heat reuse, whether internally or out to the community. We must resolve the inefficiency at the front end for the industry to become a major energy supplier to the wider community.

Over the past year, we at Iceotope, have been surprised by the level of interest in liquid cooling from across the industry, whether it is legacy sites looking to introduce a point solution for HPC systems, or modern colo sites that are targeting customers with AI-led requirements that need GPU rich compute environments. Modifications to the data centre and its infrastructure may at first, be challenging. However, the benefits can include fast ROI, and a major reduction of infrastructure energy use which improves PUE. Liquid cooling is gaining converts who realise that the technology also allows for a reduction in the overall data centre footprint, in line with predicted future smaller site requirements.

So, how do these front-end benefits assist in developing district heating? The infrastructure required for liquid cooling ensures that heat generation at the server is extracted at a much higher temperature. This is down to the dielectric coolants being around 1000 times more effective than air. This means when the coolant/water loop heat exchange takes place the water temperature is around 50oC, which can then be used in various ways, either to heat the data centre offices, or externally through an insulated pipe network.



There are numerous applications and companies that specialise in designing and implementing the network infrastructure for low carbon, recycled heat into the community.

Today, data centres should view district heating projects as part of the organisation's CSR commitment, as well as a new revenue stream and an opportunity to achieve net zero carbon status. Moreover, liquid cooling is influencing changes in data centre design and creating far more efficient and sustainable sites. These will become even more important to their local communities, as they start to heat schools, colleges, and community swimming pools.

Rich Root

Digital Marketing Specialist, ATTO Technology

atto.com



Unleash the power of flash with fibre channel


THE EVOLUTION of storage infrastructure has reached a new inflection point with the maturation of flash technology. Flash-based storage is now commonplace in server architectures, joining spinning disk and tape in secondary storage environments. Flash technology, used in conjunction with big data applications, may eventually push secondary storage architectures completely away from disk.

As the next-generation data centre emerges, where flash arrays support high-performance enterprise applications, the connection to these arrays is a critical architectural component to ensure optimal performance. By utilising a purpose-built, lossless fabric, a Fibre Channel-based storage area network (SAN) is still the best way to unleash the power of flash storage and deliver the highest performance for mission-critical, latency-sensitive applications.

Fibre channel drives better performance from flash storage

All-Flash Arrays (AFAs) speed access to data, doubling IOPs while reducing latency. However, this performance can only be realised if the connection between high-power workstations, enterprise servers and storage is matched for speed, reliability and consistency in data transfers. The best way to achieve this is by upgrading a legacy Fibre Channel SAN to the latest 32Gb Gen 7 technology.

A Fibre Channel SAN is a purpose-built technology that uses a block-level protocol for speeding access to storage and is already deployed widely in most enterprises. Unlike Ethernet, Fibre Channel is highly efficient and has very low overhead and low latency. Data delivery is fast, in-order, lossless and predictable. With the advent of Gen 7 Fibre Channel, the latest 32Gb Fibre Channel cards are capable of up to 12.8GB/s and 5 million IOPs, which sustains performance to a considerably



greater number of flash drives. Indeed, it seems as if Fibre Channel was built specifically for low-latency flash storage with 48% less overhead per frame compared with Ethernet and the ability to handle a 40% bigger workload without requiring re-transmission or error checking.

The value of tuned technology

High-performance workloads like content creation rely on high bit rate and bit depth applications that can place heavy demands on storage networks.

Latency is a significant cause of glitches, dropped data, and other network gremlins. The higher the latency the greater the delay before transferring data. This leads to poor performance and inefficiencies, and is a characteristic of networks that affects flash storage just as much as any other storage technology.

One way to address these challenges is with technology that leverages failover, load balancing protection and latency management to ensure consistent uninterrupted access to data.

Future outlook

Flash-based arrays are well on their way to dominating primary storage environments. Many factors including the decreasing overall cost per Gb of flash compared to HDD, will allow solution architects to craft an all-flash or hybrid solution depending on cost considerations and business objectives.

As latency continues to be a key performance factor for real-time applications, the full benefits of flash storage might only be realised by a tuned storage network made with Fibre Channel technology.

PHYSICAL SECURITY

Fencing, access control systems, biometrics, fire suppression, compliance and regulations – not, perhaps, as exciting as AI, cloud and the edge, but a crucial part of the overall data centre operational landscape.



IN ASSOCIATION WITH



Paul Speciale
Chief Product Officer, Scality
scality.com

Is object storage the new primary storage?

A RECENT INDEPENDENT SURVEY found that 90% of European decision-makers believe flash-based object storage will replace primary storage. And research by ESG found that 77% of IT pros who rely on all-flash object storage say the technology has had a high impact on, or has been a game-changer in their on-premises storage infrastructure. So what is behind this transformation?

The demands on IT have changed dramatically in recent years, and we see new dynamics that further push the envelope. New users, applications and growth of data volumes are some of the key drivers. These are fueled by new cloud-native and massively data-centric workloads in AI/ML, analytics and a wide range of media-based cloud services.

As businesses become increasingly digital and dependent on rapid deployment of these new applications, IT will have to build on the agility and on-demand flexibility of cloud computing. Software-based datacentre infrastructure has become the default through server virtualisation and private-cloud platforms, to software-defined networking and software-defined storage. IT must become fully proficient in these cornerstone technologies, and prepare for the next wave of cloud-native applications deployed on Kubernetes.

To support the petabyte-scale data storage and management requirements, many companies have already turned to object storage as a scalable and cost-efficient way to manage their unstructured data. We now see the next inflection point in opening up a broader set of use cases to object

storage through the use of all-flash storage. New generations of flash such as QLC (Quad Level Cell) are now at a price tipping point, making it possible to leverage flash for not only high-performance, but also high-capacity storage, and at a new affordable price point.

By delivering extremely high performance, all-flash object storage accelerates AI/ML environments, empowers business intelligence/analytics, and increases infrastructure performance and utilisation. At the same time, it lowers TCO and it reduces hardware expenditure and operational expenses. Performance of all-flash object storage can now compete with traditional all-flash block and file arrays, to provide extreme levels of throughput and sub-millisecond latencies. This opens the door to these performance intensive workloads to utilise object storage as their primary storage tier.

ML applications specifically can leverage all-flash object storage to gain high-throughput access to massive volumes of data, demanding a combination of speed and scalability from storage systems not seen before. New generations of analytics

applications consume mountains of data to provide meaningful business insights. These depend on rapid scalable storage. And edge applications for sensors, IoT devices, cameras, autonomous vehicles and thousands of others will demand the same high levels of scale and speed from their storage.

Across all of these areas, applications have adopted the AWS S3 API as a standard storage protocol for leveraging object storage. This will continue, as new cloud-native applications also naturally use these APIs and will create greater demand for all-flash object storage.





STORAGE



BACKUP
& RECOVERY



BUSINESS
CONTINUITY

Get the easiest to use tool and **ensure your data reliability and integrity** with optimal data **storage, protection** and **recovery**.

What you can expect from the Open-E JovianDSS?

- Safe and Fast Hyper-converged Storage
- On and Off-site Data Protection
- Unlimited Scalability
- High Availability Clustering

Start the free trial



Darren Watkins

Managing Director, VIRTUS Data Centres

virtusdatacentres.com



Seven layers of physical security for data centres

IF YOU'VE EVER TRIED to get into a data centre without prior permission or the correct form of identification, you'll know how stringent security is. As the keepers of, arguably, one of the world's most valuable assets – data – data centre providers have no bounds when it comes to safeguarding their customers' priceless infrastructure; the loss of data or system shut-down potentially can cost businesses millions.

Data centre security is about minimising risk and maximising operational uptime. Every data centre should be designed, built and maintained to withstand everything from corporate espionage, to terrorists, to natural disasters, to thieves.

Entry to data centres should be tightly managed with strict procedures to monitor and control visitor access both into and within different areas of the data centre. Not only is the physical security stopping criminals getting in, it is also there to delay their chances of a successful breach.

To achieve gold standard security, there should be seven layers of physical security:

1. A physical barrier

A fence that is a minimum of three metres high (five metres in some places, depending on who or what is located next door)

2. Trembler wire

A wire on top of the fence that will set off an alarm if anyone kicks, climbs or jumps over it

3. Surveillance cameras

CCTV around the perimeter of the building at all entrances and exits as well as at every access point throughout the building. A combination of motion-detection devices, low-light cameras, pan-tilt-zoom cameras and standard fixed cameras is ideal. Footage should be digitally recorded and stored offsite

4. 24/7 security guards

Always have more than one guard – one to man the systems and one or more to do a regular walk around to check the perimeter and the various data halls and plant rooms

5. Vehicle trap

Access to the facility compound, usually a parking area, needs to be strictly controlled either with a gated entry that can be opened remotely by reception or security once the driver has been identified, or with retractable bollards. This measure is to not only prevent unauthorised visitors from driving into the parking area and having a look around, but also to prevent anyone from coming straight into this area with the intention of ramming the building for access

6. Full authentication & access policy control

To get inside, government issued photo ID should be required. Once inspected, they should be given a formal ID card that allows them into different parts of the data centre, depending on whether they are a customer or a visitor. If they are visitors, they should always be accompanied

7. Biometrics

To gain access to the buildings, data floors and individual areas, biometrics should be used as a form of identification to ensure secure, single-person entry

Like cyber security, the demands on physical security are increasing as businesses recognise how much of their operations outside of the traditional IT functions are becoming dependent on the data centre. When it comes to the physical security of your data centre, a proven track record and solid experience really matters. Experienced colocation providers will offer a full suite of benefits, including service level guarantees that offer 100 per cent availability of power and cooling, 24/7 physical security and compliance with all the latest ISO certifications for managing critical infrastructure.



Jan Karremans

Director of Product Marketing, EMEA, EDB
enterprisedb.com

Database security is of utmost importance - and open source is the answer

2020's sudden shift to remote work presented organisations with a major challenge: to protect a skyrocketing volume of data from vulnerabilities. With many countries beginning to embrace a 'new normal', organisations are confronted with a new challenge - to protect the data of a hybrid workforce, as well as the data of their stakeholders.

The answer lies deeper than database security - and requires the entire data centre. In the past, the focus has been on code vulnerabilities and access but, in an era of microservices, distributed systems and online data processing, the focus will be on securing communications channels - especially due to the 'new normal' of work.

Expect to see an increased focus on interprocess security and more secure communications between applications. We must learn how to trust the information from the source, ensuring this information goes to the right people and the data is safe and properly authenticated and authorised.

The challenge here lies in the fact that many organisations use databases that were built over 30 years ago. Developing a new version of these databases, in closed source settings, most often means the existing versions must also continue working - leaving the product with a technical mortgage.

This can leave data exposed, opening up fertile ground for bugs and security leaks. If 30-year-old technology doesn't evolve, you end up with a clunky, muddled system. This is where open source database technologies like Postgres succeed. The community is passionate about building for strength, and has no fear of making changes to things that no longer work as well as they did in the past.

While misconceptions remain about open source technology, the challenges of closed source databases must also be considered. For instance, customers are often required to spend more to maintain these databases than an open source alternative. The drive of open source is to ensure organisations are not held back by features that won't repay their time in value. Open source allows them to secure the databases to



IN ASSOCIATION WITH






make sure that clusters are hyper personalised - ensuring the entire database isn't opened up to vulnerabilities. A closed database is perceived to be more secure because it's backed by a single vendor that has a strong focus on security. If you want to exploit a closed source database, you may struggle to find a vulnerability - and the vendor is often working hard to patch any holes and maintain its security.

However, if anyone can expose open source vulnerabilities, then, by the same token, anyone can identify and fix them - allowing prevention measures to be built quicker and more securely without relying on an individual to safeguard them. As more people are reading the source code, more people are able to build patches and fixes - making the core code safer. While it can be argued that open source databases lack a

specific vendor, commercial partners which provide enterprise-ready software and services can bridge the difference. Organisations like EDB can support the open source community by helping to implement safeguarding, and picking up the gaps that may remain - ensuring that patches are released early and as often as possible, whilst providing a 24/7 service that helps organisations remain 'always on', protecting themselves and their stakeholders.

As the transition to wide-scale hybrid working begins, we're not where we need to be yet - there's still much to be done to properly secure data in the 'new normal' environment. However, organisations can best protect themselves and their stakeholders by investing in open source databases to ensure peace of mind - at a fair cost.

While misconceptions remain about open source technology, the challenges of closed source databases must also be considered. For instance, customers are often required to spend more to maintain these databases than an open source alternative. The drive of open source is to ensure organisations are not held back by features that won't repay their time in value. Open source allows them to secure the databases to make sure that clusters are hyper personalised - ensuring the entire database isn't opened up to vulnerabilities



**No matter the environment,
Starline's at the centre.**

Hyperscale, Colocation, Enterprise.
Time-tested power distribution for every environment.

Starline Track Busway has been the leading overhead power distribution provider—and a critical infrastructure component—for all types of data centres over the past 30 years. The system requires little to no maintenance, and has earned a reputation of reliability due to its innovative busbar design.

Starline[®]
A brand of **legrand**

StarlineDataCentre.com/DCS

Unit C Island Road | Reading RG2 0RP | UK | +44 (0) 1183 043180