DCD
February/March 2018
datacenterdynamics.com

Innovation on ICE
Sweden's bid to make us all more efficient

Who owns DCIM data?
Cloud services create new problems

After Red Storm
The supercomputer that shot down a satellite

One megawatt.
One minute.
Zero batteries.

PILLER
Power Systems

DRIVEN BY MOTION
ACTIVEPOWER
A Division of Piller Power Systems Inc.

The battery-free data centre is here: Introducing the **POWERBRIDGE™ PB60+**, the world's largest, commercially available kinetic energy storage device for UPS applications. Delivering power above 3MW and providing 1MW of electrical power for over 60 seconds, the **POWERBRIDGE™ PB60+** achieves all this and more:

- **Reduces maintenance**
- **Eliminates replacement batteries**
- **Improves reliability**

When everything rides on reputation, now is the time to ensure your power and data are reliably protected.

Nothing protects quite like Piller

ROTARY UPS SYSTEMS / HYBRID ROTARY UPS SYSTEMS / DIESEL ROTARY UPS SYSTEMS / STATIC UPS SYSTEMS / STATIC TRANSFER SWITCHES
KINETIC ENERGY STORAGE / AIRCRAFT GROUND POWER SYSTEMS / FREQUENCY CONVERTERS / NAVAL POWER SUPPLIES / SYSTEM INTEGRATION

A Langley Holdings Company

HEADQUARTERS: PILLER GROUP GmbH
Piller France / Piller Germany / Piller Iberica / Piller Italy / Piller UK / Piller Brazil / Piller Canada / Piller USA / Piller Australia / Piller China
Piller India / Piller Korea / Piller Singapore / Piller Taiwan

ISSN 2058-4946

Contents

February/March 2018

EDITOR'S PICK

After the Storm
Charting the history of Red Storm, a supercomputer that saved Cray, gave birth to a generation of HPC systems and shot down a rogue satellite

From the Editor

Taking control with demand reduction

Data centers consume a lot of power and we want to run them more efficiently. Meanwhile, the rest of the world is increasingly worried about the reliability and sustainability of its supplies.

Data center people can help with both these issues, by taking control of what they use and how.

Demand reduction programs (p16) are about facilities cutting their demands to help grid resources. They can feed energy back into the grid, or they can simply opt to run the data center on batteries for specified periods when demand is high, effectively erasing the site from the grid.

In March, we head to Stockholm for Energy Smart - an event about efficient infrastructure

Our Energy Smart event in Stockholm in March will have demand reduction on the agenda - along with many other things - when we bring together power people and facilities folks, in one of the world's most energy-conscious countries (p14).

Kicked off by the State Secretary of Sweden's Ministry for Enterprise and Innovation, the event will explore district heating and fuel-cell powered facilities (p21), and look at the energy economics of cryptocurrency mining.

Come and join us - and you will leave empowered!

You don't have to be in Sweden's cool climate to operate efficiently. On p39 we look at a data center which achieved an enviable PUE in Africa.

We also look at a landmark machine from HPC history (p30). Red Storm saved Cray, revolutionized supercomputing - shook the leadership of Intel in the field - and shot down a dangerous satellite.

Manufacturing will be making some of the biggest demands on digital infrastructure. As it moves to "Industry 4.0," the level of automation will increase ever more rapidly, binding networked resources closer together and producing increased demands on the resources to handle that data.

It boils down to a set of demands that combines the Internet of Things, edge networking and an absolute need for real-time response. This makes for a unique and demanding set of requirements (p25).

The air in data centers is as important as they power they consume, and data centers are suffering surprising and unintended consequences from two different green initiatives (p28).

Regulations to remove lead from electronic systems leave circuit boards vulnerable to corrosion, while free cooling increases contamination risks.

How do you deal with that? A piece in this magazine - abridged from a longer article on the DCD website - gives you strategies.

Take control. It's the key to staying reliable and becoming more efficient.

bit.ly/DCDmagazine

$3.3trn

Total market cap value of the top five public firms: Apple, Alphabet (Google), Microsoft, Amazon and Facebook (Dec 17). If they were a country it would be the world's fifth largest, just behind Germany's GDP

Peter Judge
DCD Global Editor

Subscriptions
datacenterdynamics.com/magazine

To email one of our team
firstname.surname@datacenterdynamics.com

Find us online
datacenterdynamics.com | dcd.events | dcdwards.global | dcpro.training

DatacenterDynamics
DCDnews
DatacenterDynamics

Meet the team

Global Editor
Peter Judge
@Judgecorp

News Editor
Max Smolaks
@MaxSmolax

Reporter
Sebastian Moss
@SebMoss

Reporter
Tanwen Dawn-Hiscox
@Tanwendh

Editor LATAM
Virginia Toledo
@DCDNoticias

Assistant Editor LATAM
Celia Villarrubia
@DCDNoticias

SEA Correspondent
Paul Mah
@PaulMah

Brazil Correspondent
Tatiane Aquim
@DCDFocuspt

Head of Design
Chris Perrins

Designer
Mar Pérez

Designer
Ellie James

Head of Sales
Yash Puwar

Global Account Manager
Aiden Powell

Global Conference Director
Rebecca Davison

Conference Director
Giovanni Zappulo

Conference Manager, EMEA
Merima Dzanic

Head office

DatacenterDynamics
102–108 Clifton Street
London EC2A 4HW
+44 (0) 207 377 1907

Narada®

High Rate, Long Life Batteries for Critical Power - Data Centers

RELIABLE

When you need long term reliable backup power, Narada High Rate and High Operating Temperature batteries deliver. Engineered and manufactured in state of the art facilities, Narada provides solutions to meet your Critical Power needs.

ISO9001/14001/TL9000 Certified Quality
Backed by Industry Leading Warranties

Narada...Reliable Battery Solutions

www.mpinarada.com - email : ups@mpinarada.com - MPI Narada - Newton, MA
Tel: 800-982-4339

DCPRO Mission Critical Training

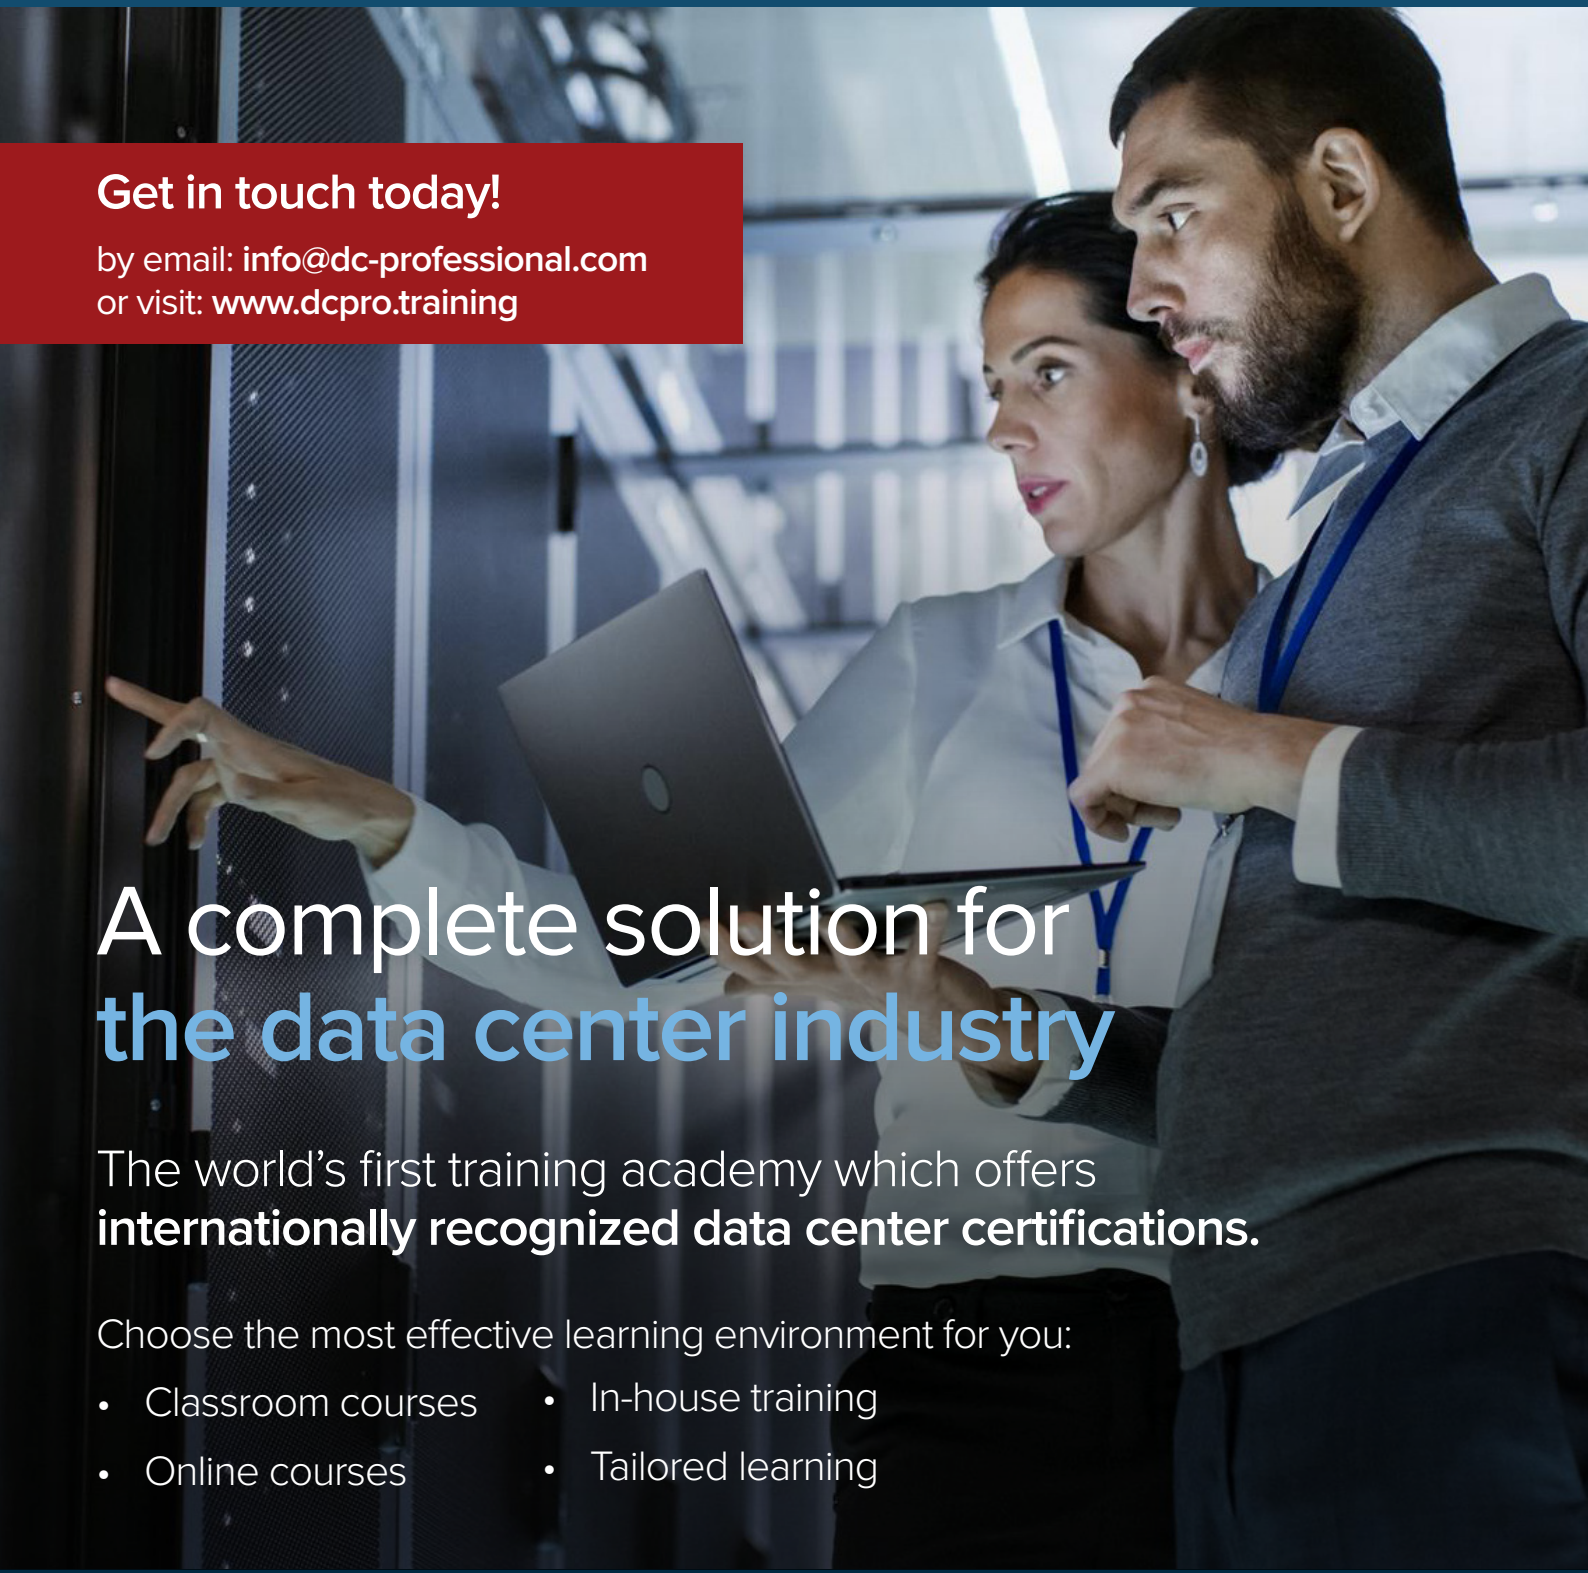

Get in touch today!

by email: **info@dc-professional.com**
or visit: **www.dcpro.training**

A complete solution for the data center industry

The world's first training academy which offers **internationally recognized data center certifications.**

Choose the most effective learning environment for you:

- Classroom courses
- Online courses
- In-house training
- Tailored learning

White space

A world connected: The biggest data center news stories of the last two months

Apple to spend more than $10bn on US data centers over five years

Technology giant Apple has promised to increase its investment in the United States over the next five years, including more than $10 billion spent on data centers.

The "expanded capital expenditure" on data centers adds to what the company had already planned to spend on its digital infrastructure.

Overall, Apple plans to invest more than $30bn in its US operations, establishing a new campus and creating more than 20,000 jobs. Taking advantage of a temporary tax discount on cash repatriation, the company will bring most of its $250bn hoard back to the US, paying a one-off tax bill of roughly $38bn - which it said is the largest tax payment of its kind in history.

The news was announced as Apple CEO Tim Cook visited a groundbreaking ceremony for the company's data center expansion in Reno in January.

In addition to the Reno facility, Apple operates its own data center sites in seven US states, including North Carolina, Oregon, Nevada, Arizona and Iowa.

The company reportedly uses Google Cloud Platform, Microsoft Azure and Amazon Web Services, but is thought to be trying to reduce its reliance on cloud goliaths that compete with it elsewhere.

"Apple is a success story that could only have happened in America, and we are proud to build on our long history of support for the US economy," Cook said.

"We believe deeply in the power of American ingenuity, and we are focusing our investments in areas where we can have a direct impact on job creation and job preparedness. We have a deep sense of responsibility to give back to our country and the people who help make our success possible."

The company will also increase the size of its Advanced Manufacturing Fund, which is meant to support US manufacturers, from $1bn to $5bn. In addition, it will give a $2,500 stock grant to a large number of its employees.

bit.ly/AsAmericanAsApplePie

News in brief

Equinix to acquire 1.6m sq ft Infomart Dallas for $800m
This leaves ASB Real Estate Investments' Infomart Data Centers with three US facilities and adds a key hub to Equinix's portfolio.

Apple to build second Chinese data center
Tech giant Apple will build a $160m data center in Inner Mongolia, China, joining its under-construction facility in the southwestern Guizhou Province.

Statkraft seeks partners for large data center projects in Norway
The state-owned Norwegian power company is in talks with potential partners as it looks to build large data centers in Norway, after recent tax changes.

Tencent opens data center in Chongqing, China
The new facility aims to target the manufacturing sector - with Chongqing a major production hub, and the economic center of the upstream Yangtze basin.

Google invests €250m in its third data center in Belgium
Its first project to feature its own solar farm.

Data center mergers, acquisitions totaled $20bn in 2017

Major merger and acquisition transactions involving data center operators reached a total of $20 billion in 2017, according to Synergy Research. The number is higher than the total for both 2015 and 2016 combined, with Digital Realty and Equinix the year's biggest spenders.

Vertiv buys PDU maker Geist

American power and cooling specialist Vertiv has acquired Geist, the manufacturer of power distribution units (PDUs) and associated data center products. Financial terms of the transaction were not disclosed.

The announcement came a week after Vertiv closed the acquisition of thermal management company Energy Labs.

Geist is headquartered in Lincoln, Nebraska, and has more than 250 employees. Following the acquisition, the company's president, Brad Wilson, will remain manager of the business.

Geist's data center power, cooling, monitoring, and infrastructure management products will complement Vertiv's existing PDU line-up.

When asked whether this was a step to consolidate the PDU market, Vertiv's VP of business development, Jack Pouchet, told *DCD* that "there is still room for growth and innovation in this sector."

bit.ly/GoingGoingGeist

Meltdown and Spectre CPU vulnerabilities uncovered, patches do more harm than good

Researchers uncovered speculative execution vulnerabilities in the majority of Intel's CPUs from the last 15 years, as well as some AMD, Arm and IBM chips.

Known as Meltdown and Spectre, the issues were kept secret for several months, before being made public in January. Then the real problems started.

A software patch offered by Intel increased the likelihood of server reboots. The company secretly told select customers not to install the patches, before finally sharing its concerns with the wider industry a week later. It also originally denied the fixes would impact performance, but has now admitted that it has observed slowdowns on affected systems.

"While the firmware updates are effective at mitigating exposure to the security issues, customers have reported more frequent reboots on firmware updated systems," Navin Shenoy, executive VP and GM of Intel's Data Center Group, said.

"We have reproduced these issues internally and are making progress toward identifying the root cause. In parallel, we will be providing beta microcode to vendors for validation [soon]."

The vulnerabilities were revealed too late to have any impact on Intel's Q4 bottom line, however. And, with the company's near-monopoly on x86 chip sales, it is unclear what the long-term effects will be.

But Intel did warn in its earnings release that Meltdown and Spectre could hurt future financial performance, along with customer relationships and the company's reputation.

Intel CEO Brian Krzanich said: "Security is a top priority. We will restore confidence in data security with customer-first urgency, transparency, and timely communication."

Later this year, Intel expects to release chips that address both flaws at a hardware level.

bit.ly/NoMeltdownForIntelsFinances

Led by ex-Intel execs, Ampere plans new 64-bit Arm server processors

A start-up led by the former president of Intel, Renee James, has set its sights on the data center market.

64-bit Arm-based processors from the new company, Ampere, are sampling now and are expected to be in production in the second half of the year.

Ampere's products operate at up to 3.3 GHz, with a power envelope of 125 watts.

"We have an opportunity with cloud computing to take a fresh approach with products that are built to address the new software ecosystem," James said.

"The workloads moving to the cloud require more memory, and at the same time, customers have stringent requirements for power, size and costs.

"The Ampere team's approach and architecture meets the expectation on performance and power and gives customers the freedom to accelerate the delivery of the most memory-intensive applications and workloads such as AI, big data, storage and database in their next-generation data centers."

In addition to James, the chief architect, Atiq Bajwa, spent 30 years at Intel, the EVP of hardware engineering, Rohit Vadwans, worked there for 26 years, and the CFO and COO, Chi Miller, was both an Intel and Apple executive. Senior fellow Greg Favor comes from AMD.

Ampere is based in Santa Clara, California, with its headquarters just a four-minute drive from Intel's offices.

bit.ly/AmpingUp

Peter's chip factoid

Ampere is the SI unit of electric current, named after scientist André-Marie Ampère. Ampere is also the name of a French company which made cars with an electric clutch from 1906 to 1909

Vox Box

Ghada Badawy & Souvik Pal
Researchers
CIRC McMaster University
Tell us about your data center research projects
Badawy: We're building wireless sensor nodes. We're not just collecting the information - we've built some machine learning algorithms that will predict what's going to happen next.
Pal: We're working on a whole new architecture which can deliver cooling at the rack level. A cooling unit that looks like a 2U server, that you can mount in any 2U space.

bit.ly/DCDSmartSensors

Rebe Pronovost
Director of operations
Maya HTT
What have you found from adding AI features to the Datacenter Clarity DCIM product?
We've been doing AI projects for years, and from our experience it all comes down to the data. In AI projects the challenge comes down to getting good quality data. You often hear that good data with a bad AI is better than bad data with a good AI.
It all comes down to having the right data. That all comes down to its labeling, its organization and its classification.

bit.ly/DCDGoodData

DCD > Debates

A new webinar format bringing the dynamism of our live panel discussions to a global audience

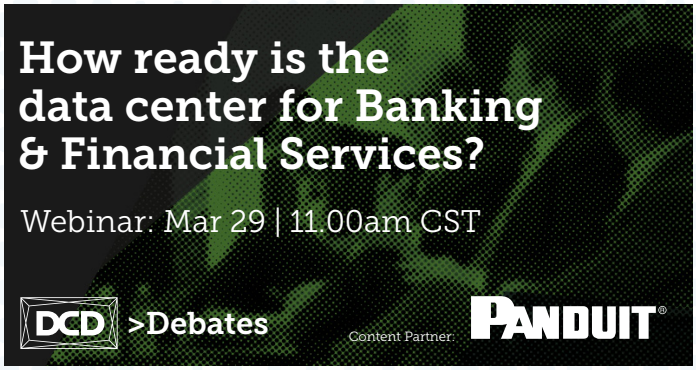

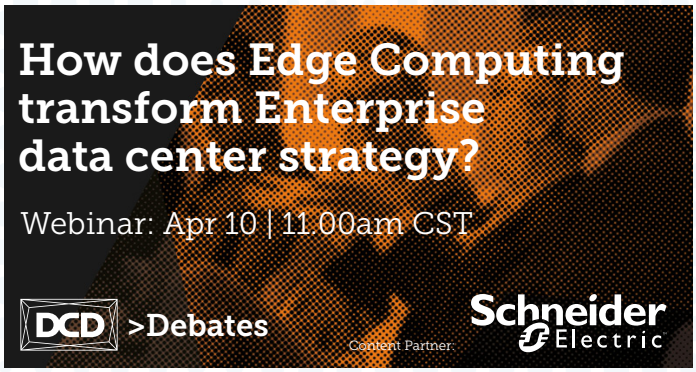

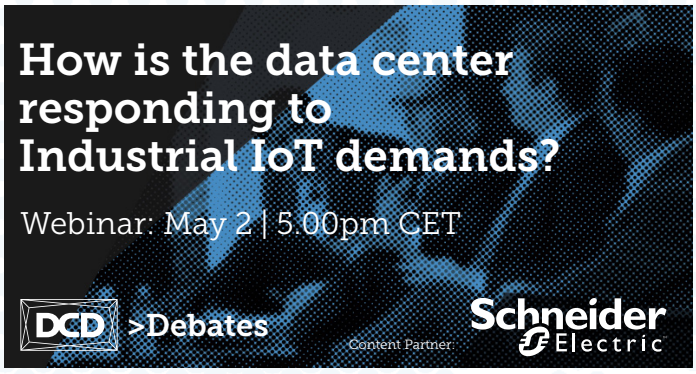

Coming soon ...

Is Machine Learning and AI ready for mainstream in the data center?
Webinar: May 16 | 11.00am CST - 4.00pm GMT

Will smart energy tech unite the data center and energy industries?
Webinar: May 22 | 11.00am CST - 4.00pm GMT

How is HSDT tech solving hyperscale data center congestion?
Webinar: Jun 6 | 11.00am CST

Expert Panel: What's new in data center cooling?
Webinar: Jun 13 | 11.00am CST - 4.00pm GMT

Is colo ready for crypto-mining and the blockchain?
Webinar: May 29 | 11.00am CST

Resilient Infrastructure: Is the future of disaster recovery in the cloud?
Webinar: Jun 12 | 11.00am CST - 4.00pm GMT

Successful transitions from on-prem to colo; what's best practice?
Webinar: Jun 19 | 11.00am CST - 4.00pm GMT

Find out more...

dcddebates.com

In the third quarter of 2017, more than half of all colo revenue came from just 20 metros

Of all the revenue generated by retail and wholesale colocation around the world in the third quarter of 2017, 59 percent came from only 20 metropolitan areas, according to a report by Synergy Research Group.

The top five were Washington DC, New York, Tokyo, London and Shanghai, which together accounted for 26 percent of total revenue, followed by Chicago, Dallas, Frankfurt, San Francisco and Singapore, which accounted for 15 percent of revenue.

Finally, Amsterdam, Atlanta, Beijing, Hong Kong, Los Angeles, Paris, Phoenix, Seattle, Sydney and Toronto generated 18 percent of revenue, and the rest of the world accounted for the remaining 41 percent.

Reinforcing this trend, revenue growth in the top five metropolitan areas year-on-year was reportedly two percent higher than in the rest of the world, indicating an ongoing concentration of the market.

The Washington DC and North Virginia metropolitan area stood apart from the rest, emerging as "by far the largest wholesale market in the world," according to John Dinsdale, Synergy's chief analyst and research director. This market reported 20 percent growth in the last four quarters.

bit.ly/ConcentratedintheHandsoftheFew

Equinix CEO resigns over "poor judgment" in employee matter

Steve Smith, the CEO of global interconnection and data center company Equinix, has resigned with immediate effect "after exercising poor judgment with respect to an employee matter."

The Equinix Board of Directors, "in the best interests of the company," accepted the resignation and has appointed executive chairman Peter Van Camp as interim CEO.

Equinix provided no further details about the reason for the decision.

"The Board gave this matter the deepest consideration and recognizes the many contributions Steve made over the past 11 years to achieve the global scale, reach and market leadership the company enjoys today," Van Camp said.

"He has worked hard to grow and sustain the business, and we greatly appreciate his efforts. I also want to emphasize that this action was not related to the company's operational performance or financial condition, both of which remain strong. The Board and leadership team remain fully committed to the strategy."

During his time as CEO of Equinix, Smith, 61, grew the company's annual sales from $400 million to a projected $4.35 billion for 2017.

In his most recent company blog post, Smith said: "As I look ahead, I know we have the potential to be remembered as one of the companies that enabled the cloud to become "real" globally, or helped the IoT move from a concept to a central driver of business, or helped bring to life any other tech trend you can name. I want to be in the middle of what makes it all work.

"And I'm deeply thankful for the time I have already spent at the center of all that opportunity, helping people and companies connect to it, here at Equinix."

Smith's replacement, Van Camp, has worked at Equinix for more than 17 years, including serving as CEO from 2000 to 2007. His new position could be temporary, however, with the board commencing a formal process to appoint a new CEO.

"The company is well-positioned strategically, with tremendous depth at the leadership level and a passionate team that will guide the business and continue to drive the performance of the company," Van Camp added.

bit.ly/ExitingEquinix

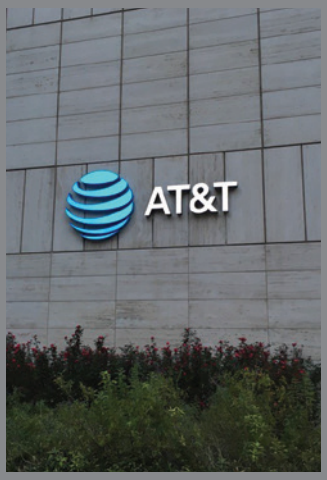

AT&T considers sale of its data center business

AT&T is once again considering a sale of its data center business, exploring a divestiture of its facilities, reports *The Wall Street Journal*.

The telecoms giant looked into a similar sale for $2 billion back in early 2015, but eventually settled on selling its hosting business to IBM.

The process is still ongoing and, like before, may not result in a sale. The business is thought to generate about $135 million in earnings before interest, taxes, depreciation and amortization, the sources said.

Should the deal go ahead, it will follow in the footsteps of several other US telcos divesting their data center assets. Last year, Verizon sold its data centers to Equinix for $3.6bn, CenturyLink sold its facilities to a BC Partners-led consortium for $2.3bn, and, in 2015, Windstream sold its facilities to TierPoint for $575m.

AT&T has already taken steps to move to the cloud. In 2016, it announced a multi-year "strategic relationship" with AWS to link to business customers, and said that it would run its applications on IBM's cloud.

bit.ly/SecondTimeLucky

Eni launches world's most powerful industrial supercomputer

Italian multinational oil and gas giant Eni has launched a high performance computing (HPC) system that outperforms any other industrial supercomputer.

The HPC4, located at the Eni Green Data Center in Italy, delivers peak performance of 18.6 petaflops. When combined with the existing HPC3, the system reaches a computational peak capacity of 22.4 petaflops.

HPC4 is the only non-governmental and non-institutional system ranked among the top ten most powerful systems in the world, when checked against the latest Top500 supercomputer list from last November.

Built out of hardware supplied by HPE, HPC4 is based on 16,000 ProLiant DL380 nodes, each equipped with two 24-core Intel Skylake processors - totaling more than 76,000 cores - and two Nvidia Tesla P100 GPUs, connected through a high-speed Enhanced Data Rate InfiniBand fabric.

The machine is linked to a high performance 15-petabyte storage subsystem.

The data center reportedly has a power usage effectiveness (PUE) of 1.2 and uses free air cooling to help lower operational costs.

"The investments devoted to reinforcing the supercomputing infrastructure and the development of algorithms are a significant part of Eni's digital transformation process," Claudio Descalzi, CEO of Eni, said.

"We can store and process enormous quantities of data for geophysical imaging, the modeling of oil systems and reservoirs, in addition to using predictive and cognitive computing algorithms for all our business activities."

bit.ly/OilPower

European Union plans €1bn supercomputer push

The European Union plans to spend €1 billion (US$1.2bn) on supercomputing by 2020, with half the funds coming from member states and the other half from the European Commission.

A new legal and funding structure, the EuroHPC Joint Undertaking, will acquire, build and deploy high performance computing infrastructure across Europe.

"Supercomputers are the engine to power the digital economy. It is a tough race and today the EU is lagging behind: we do not have any supercomputers in the world's top-ten," Andrus Ansip, European Commission VP for the Digital Single Market, said.

"With the EuroHPC initiative we want to give European researchers and companies world-leading supercomputer capacity by 2020 – to develop technologies such as artificial intelligence and build the future's everyday applications in areas like health, security or engineering."

Mariya Gabriel, Commissioner for Digital Economy and Society, added: "Supercomputers are already at the core of major advancements and innovations in many areas directly affecting the daily lives of European citizens. They can help us to develop personalized medicine, save energy and fight against climate change more efficiently.

"A better European supercomputing infrastructure holds great potential for job creation and is a key factor for the digitization of industry and increasing the competitiveness of the European economy."

bit.ly/BuildingaFutureTogether

Research: US gov't agencies are nowhere close to meeting DCOI targets

American government agencies continue to struggle with efforts to reduce their data center footprint, as required by the terms of the Data Center Optimization Initiative (DCOI).

In a survey conducted by public sector publication *MeriTalk*, fewer than 20 percent of IT managers at participating agencies said they would have been on track to meet the original DCOI deadline of September 2018.

The deadline has since been moved to 2020 – the decision was supported by 69 percent of respondents - but government bodies need to improve their efforts if they are to reach the ambitious targets of eliminating 25 percent of 'tiered' data centers and 60 percent of 'non-tiered' facilities.

DCOI is an addition to the Federal Technology Acquisition Reform Act (FITARA), set out by the Obama administration in 2014 and aimed at improving efficiency, cost-effectiveness and security of federal data centers.

The *MeriTalk* study, sponsored by ViON and Hitachi Vantara Federal, looked at the progress made in the three years since the initiative was announced. According to the results, just three in ten federal IT managers were "very satisfied" with their agency's modernization efforts to date.

At the same time, 77 percent of participants said they believed that the Modernizing Government Technology (MGT) Act, signed into law in December, would help them meet their DCOI goals.

The study found that DCOI has been quite successful in moving organizations to the cloud: 64 percent of respondents said their agency's data center acquisitions were beginning to shift towards utility-based or pay-as-you-go models.

It also found that 67 percent of IT managers who were "very satisfied" with their modernization progress had implemented DCIM solutions, compared to just 30 percent of those who were "not very satisfied."

bit.ly/DelayedEnding

Peter's consolidation factoid

The Federal Data Center Consolidation Initiative (FDCCI) of 2010 planned to close 1,200 facilities in 2010 and save up to $8bn per year. By 2014, 3,000 had been closed... and 12,000 more remained.

Uptime is everything—
So don't fall for the imitators.

Trust 30 years of innovation and reliability.

Originally released nearly 30 years ago, Starline Track Busway was the first busway of its kind and has been refining and expanding its offering ever since. The system was designed to be maintenance-free; avoiding bolted connections that require routine torqueing. In addition, Track Busway's patented u-shaped copper busbar design creates constant tension and ensures the most reliable connection to power in the industry—meaning continuous uptime for your operation. For more information visit **StarlinePower.com**

STARLINE®
TRACK BUSWAY

DCD Calendar

Stay up to date with the latest from DCD - as the global hub for all things data center related, we have everything from the latest news, to events, awards, training and research

Events

Swedish Government endorses DCD>Energy Smart

March 13 2018 // **The Brewery, Stockholm**

With green energy firmly on the national agenda, State Secretary Niklas Johansson will join us onstage at DCD>Energy Smart to discuss ways the Swedish Government is incorporating sustainable, energy efficient data centers into the national digital economic strategy in 2018 and beyond.

On March 12, delegates have the chance to visit Stockholm Exergi's KKV8 CHP facility – one of the world's largest biomass projects at the center of the fossil fuel free initiative – and Ericsson's Global ICT Center, which provides district heat to the local grid in the Stockholm area.

bit.ly/dcdenergysmart

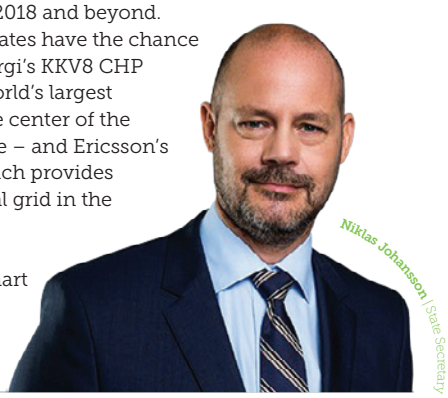
Niklas Johansson | State Secretary

Build your Edge roadmap at DCD>Enterprise

May 1-2 2018 // **Marriott Marquis, New York**

With Edge data centers set to enable the development of more robust digital architectures across the financial, retail, healthcare and manufacturing sectors, DCD>Enterprise will bring together IT professionals from across these key vertical for a highly interactive Edge pre-conference focus day on April 30.

Enabling IT professionals to design and implement their own edge roadmaps, discussions will focus on strategies for deploying data centers at the edge and an assessment of new technologies and architectures.

bit.ly/dcdenterprise-edge

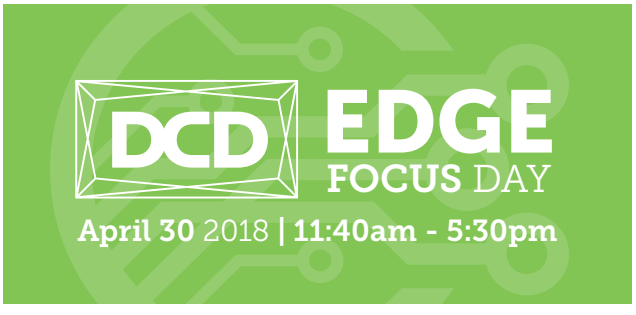

Scalability solutions – rapid growth strategies underpin the DCD>Webscale agenda

June 26 2018 // **Marriott Marquis, San Francisco**

As data center and network infrastructure executives from across the West Coast seek to rapidly scale their capabilities to meet the business demands of the zettabyte era, DCD>Webscale brings together industry leaders and emerging players in the market to discuss possible solutions, including: hyper-converged infrastructure, open source infrastructure and routes to increasing data center densities.

bit.ly/dcdwebscale

DCD>Portugal | Lisboa
Lisbon Hotel Tivoli Oriente
Apr 11 2018

DCD>Indonesia | Jakarta
Jakarta The Ritz-Carlton Jakarta, MegaKuningan
Apr 5 2018

DCD>Argentina | Buenos Aires

Apr 25 2018 // **Buenos Aires Goldencenter Eventos**

After the consolidation stage, the hybrid cloud makes its way into the Argentine data center market and users want to know which applications to migrate to the cloud, with experts on hand to share practical insights.

bit.ly/dcdargentina

DCD>Debates - A new webinar format bringing the dynamism of our live panel discussions to a global audience

With FinTech innovations threatening to disrupt operations of specialized financial data centers, join our expert panel, including Jim Fletcher, Strategy Advisor, Momenta Partners Office, and Panduit's Director of Data Center Development, Tom Kelly. Both speakers will pinpoint the ways your firm can leverage modern digital infrastructure architectures and technologies to thrive in an era of digital infrastructure transformation.

Register here: bit.ly/PanduitWebinar

DCD>Focus On | Hyderabad
Hyderabad The Westin Hyderabad Mindspace
Apr 26 2018

DCD>España | Madrid

May 23 2018 // **Madrid The N@ave**

The 11th edition of DCD>Spain, the leading event for data center and cloud infrastructure professionals in Spain returns, bringing together more than 1,000 IT, design, construction and operation executives.

bit.ly/dcdespana

DCD>Enterprise | Shanghai

Jun 7 2018 // **Shanghai Tower Conference Center**

DCD>Enterprise, the 9th annual conference and expo, is set to explore the full IT infrastructure transformation ecosystem as part of the DCD China Data Center Week. This week of events will track the world's largest emerging digital hub.

bit.ly/enterprise-china

DCD>Colombia
Ágora Bogotá Convention Center
Jun 20 2018

WASTE NOT

Max Smolaks
News Editor

Data center UPS systems can be used to support the electrical grid, creating a new revenue stream for operators

Batteries are an essential component of most data centers, where they are used as part of uninterruptible power supply (UPS) systems. Huge amounts of stored energy act as an insurance policy, enabling facilities to withstand power spikes, sags and momentary outages, and giving them a few precious seconds to switch to an alternative power source, or shut the equipment down in a controlled manner. But what if there's no outage?

Traditional lead-acid batteries have a short lifespan, take a lot of valuable space and require a carefully managed environment. They also need to be recycled by a specialist, since they are full of toxic elements. Lithium-ion batteries present an attractive alternative – they can tolerate higher temperatures and last longer, but are more expensive, and are still considered a novelty.

Most data center design guidelines – including those from the Uptime Institute – mandate the use of UPS systems to maximize reliability. But without an outage, batteries and their potential are essentially wasted.

At the same time, power grids in many parts of the world find themselves increasingly under stress, with sudden demand fluctuations leading to brownouts and threatening the conveniences of modern life. As the energy market moves away from fossil fuels, grid operators also have to deal with the unreliable nature of renewable energy sources – for example, solar panels are only productive during the day, while the performance of wind turbines depends on atmospheric conditions.

These two facts suggest an idea: why not use data center UPS systems – along with other battery clusters, like those in electric cars and smart homes

- to compensate for the difference between grid-wide generation and consumption in real-time? The idea has been discussed for several years. In this scenario, batteries store energy during the periods of peak production to later sell it back to the grid, participating in something called the Frequency Containment Reserve (FCR) market. Enabling such a market for data centers is not easy, but this approach could benefit everyone involved, from facility owners, to utilities, to consumers.

"From a hardware point of view, most modern UPS systems on the market can do this," Janne Paananen, technology manager at Eaton EMEA, told *DCD*. "It is better for the grid, and for everyone, that you maximize the use of the hardware, and at the same time, of course, you can get better value from your equipment."

Late last year, Eaton launched UPS-as-a-Reserve Service, calling it the first data center solution that lets organizations earn money from their UPS investment. The service was developed in collaboration with Fortum, one of the largest energy companies in Scandinavia. It transforms data center batteries into a virtual power plant, ready to supply energy in order to avoid grid-level outages – one of the primary reasons UPS systems are needed in the first place.

With UPS-as-a-Reserve Service, data center operators choose how much energy capacity to offer, when, and at what price. Best of all, Paananen says that this additional functionality could be carried out without any risks to the primary purpose of a modern UPS – keeping the data center online.

"Normally people understand UPS as a uni-directional device that would only take power from the grid and feed it to the load. In reality, a modern UPS with transistor-based rectifiers and inverters is a collection of bi-directional converters, it is very flexible," he said.

"When we are doing this - discharging the battery and regulating the power - instead of taking a leap of faith, opening the main breaker in the data center and hoping that everything works, we are doing it in a more controlled manner. We limit the battery discharge and it is done in parallel with the rectifier. If you have a problem with the battery, you just stop discharging and keep supporting the load.

"UPS, of course, has been designed to immediately move all the power out of batteries if you have a mains outage. Normally, if the rectifier goes off, you transfer all the power in a few milliseconds. When we support the grid, the process takes a few seconds, so it's very slow and gentle – UPS can have an afternoon tea while it's doing it."

Eaton developed the technology to enable this functionality way back in 2004, as part of an automated testing system that discharges batteries once a month. This testing routine has been used millions of times worldwide – a testament to its reliable operation – and has now found new uses in frequency regulation.

"The most important thing is to make sure you don't stress the [lead-acid] battery too much. I've read whitepapers stating that you could do this with your batteries 10-15 minutes a day – but that is creating lots of cycles with a very high depth of discharge. It would destroy the battery very quickly, and is not a smart thing to do.

"On the other hand, if you analyze the grid frequency data and you understand how the grid behaves, statistically, and you select the operating parameters so you don't stress the battery too much, going for a more dynamic type response – it doesn't impact battery life."

The lithium battery is an extra investment, but it also allows you to earn more

Lithium-ion batteries can tolerate more discharge cycles and can be kept in a partially charged state, so they are even more useful, since they can help regulate frequency either up or down. This enables participation in enhanced frequency regulation schemes where the value of the market is higher. ▶

Peak Shaving

Another way to employ data center generators and batteries for the common good is peak shaving - the process of reducing the amount of energy purchased from the utility during peak demand hours, thus saving substantial amounts of money that would have been spent on peaking charges. It also benefits utilities, helping them reduce the operational cost of generating power.

The main difference between peak shaving and schemes like UPS-as-a-Reserve Service is in the fact that the former uses on-site energy to effectively 'erase' the facility from the grid when overall power demands are exceptionally high.

Data center operators can get paid for using peak shaving techniques, but the amount of money they can earn is much lower than if they were able to actually sell energy back to the grid.

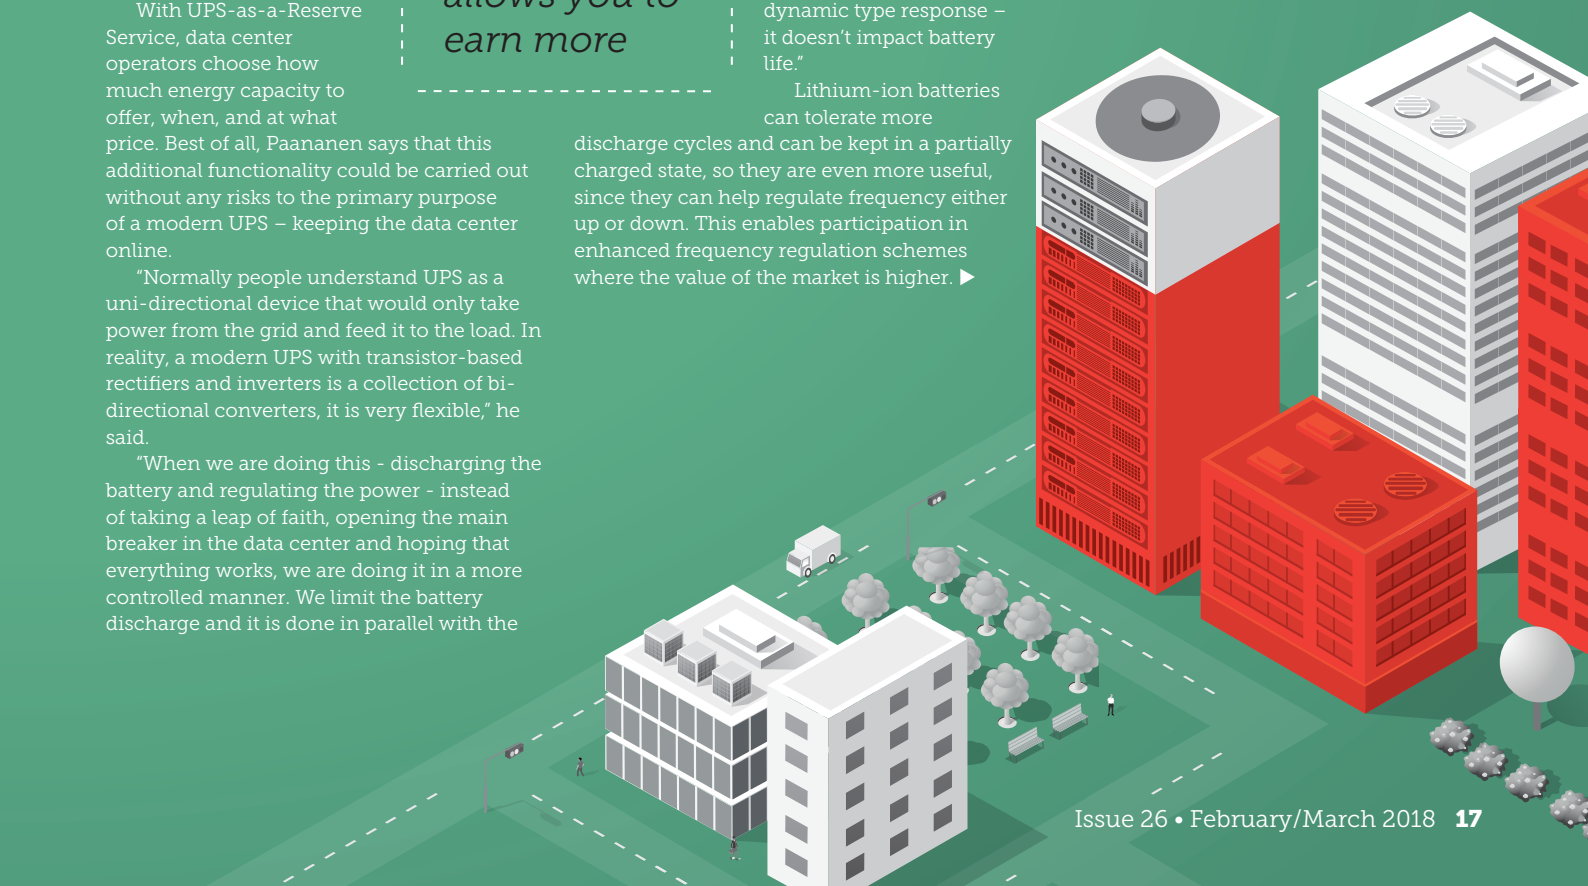

▶ "The lithium battery is an extra investment, but it also allows you to earn more," Paananen said.

Eaton's scheme has just one drawback: despite bold claims about capabilities of modern hardware, UPS-as-a-Reserve Service currently supports a limited range of UPS systems, so it can't be used in an older facility without an extensive upgrade. It also requires participation of local transmission network operators, people actually responsible for keeping the grid operational, so pilot projects are focused on Scandinavia where Eaton has made relevant arrangements with Fortum. Other compatible markets that could adopt this approach in the near future include the UK and Finland.

150MW: upcoming UPS system in South Korea

The idea of using data centers to help the grid is being investigated by several major industry players, including Schneider Electric. "We definitely see opportunities where a data center's UPS asset can be utilized beyond its primary functionality of resiliency and back-up to participate in a frequency regulation market," told us Michael Maiello, the company's SVP, technology and offer development executive.

"Since the UPS has historically been deployed for the sole intent of providing resiliency to critical loads, there is often a hesitation to utilize this asset for an alternate purpose. However, if properly managed, these concerns can be allayed by ensuring there is always sufficient energy remaining in the UPS to provide the critical backup required for the data center application."

Schneider Electric has its own UPS systems capable of this functionality.

Maiello added that the idea hasn't really caught on because of lack of understanding - the perceived risks simply outweigh potential benefits.

Paananen noted that some data center operators have already gone through the effort of plugging their generators into the grid, enabling their facilities to act like small power stations in periods of high power demand. However, he thinks this is a wrong approach to take for an industry obsessed with efficiency and sustainability: "It is a bit funny: our goal is to use renewable energy – wind power, solar power - and avoid CO2 emissions. But that makes the grid more volatile, so you need to then run generators and burn fossil fuels so that you can do green power.

"With UPS-as-a-Reserve Service, you can manage the fast changes and 99 percent of the time you would not need to start generators. You would only use them if a disturbance takes a long time."

Using UPS systems to their full potential is good for offsetting the enviromental impact of data centers, and benefits the neghborhood, but more importantly, it's great for the bottom line: according to Eaton, a data center could expect to raise up to €50,000 ($61,500) per MW of power allocated to grid support per year. For this reason alone, participation in FCR, along with waste heat recapture, should become an essential data center feature.

"In the beginning, some people will have their doubts and they will wait for other guys to try it out. But when it's done a few times, there are case studies and it's a proven concept, everybody will be doing it." ●

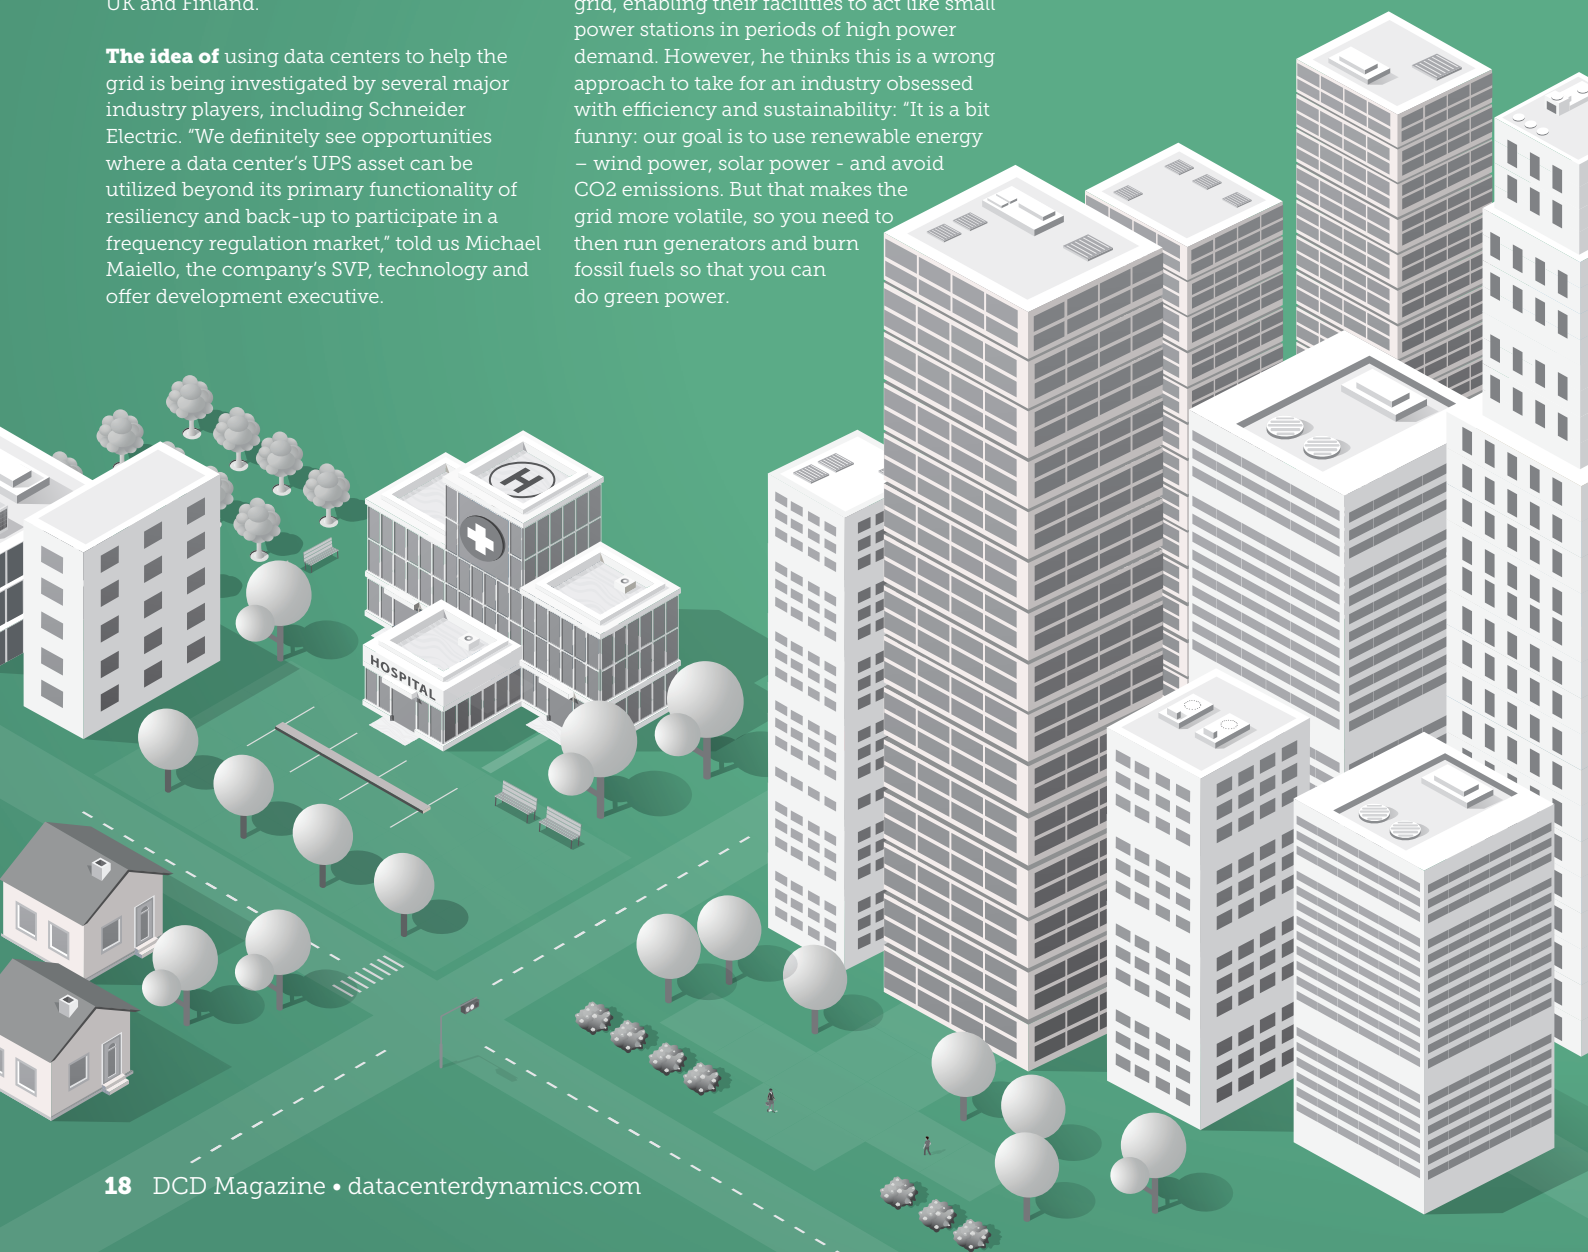

Flexible Power:
From overhead to in-rack

STARLINE

When data center power requirements call for higher density, multiple types of receptacles, and future flexibility, Starline goes above and beyond. Though it may look like other power distribution systems, Starline's innovative design provides users with the flexibility to choose and use different types of receptacles on the Cabinet Busway system. To see how Starline is changing power distribution in data centers, visit **StarlinePower.com**

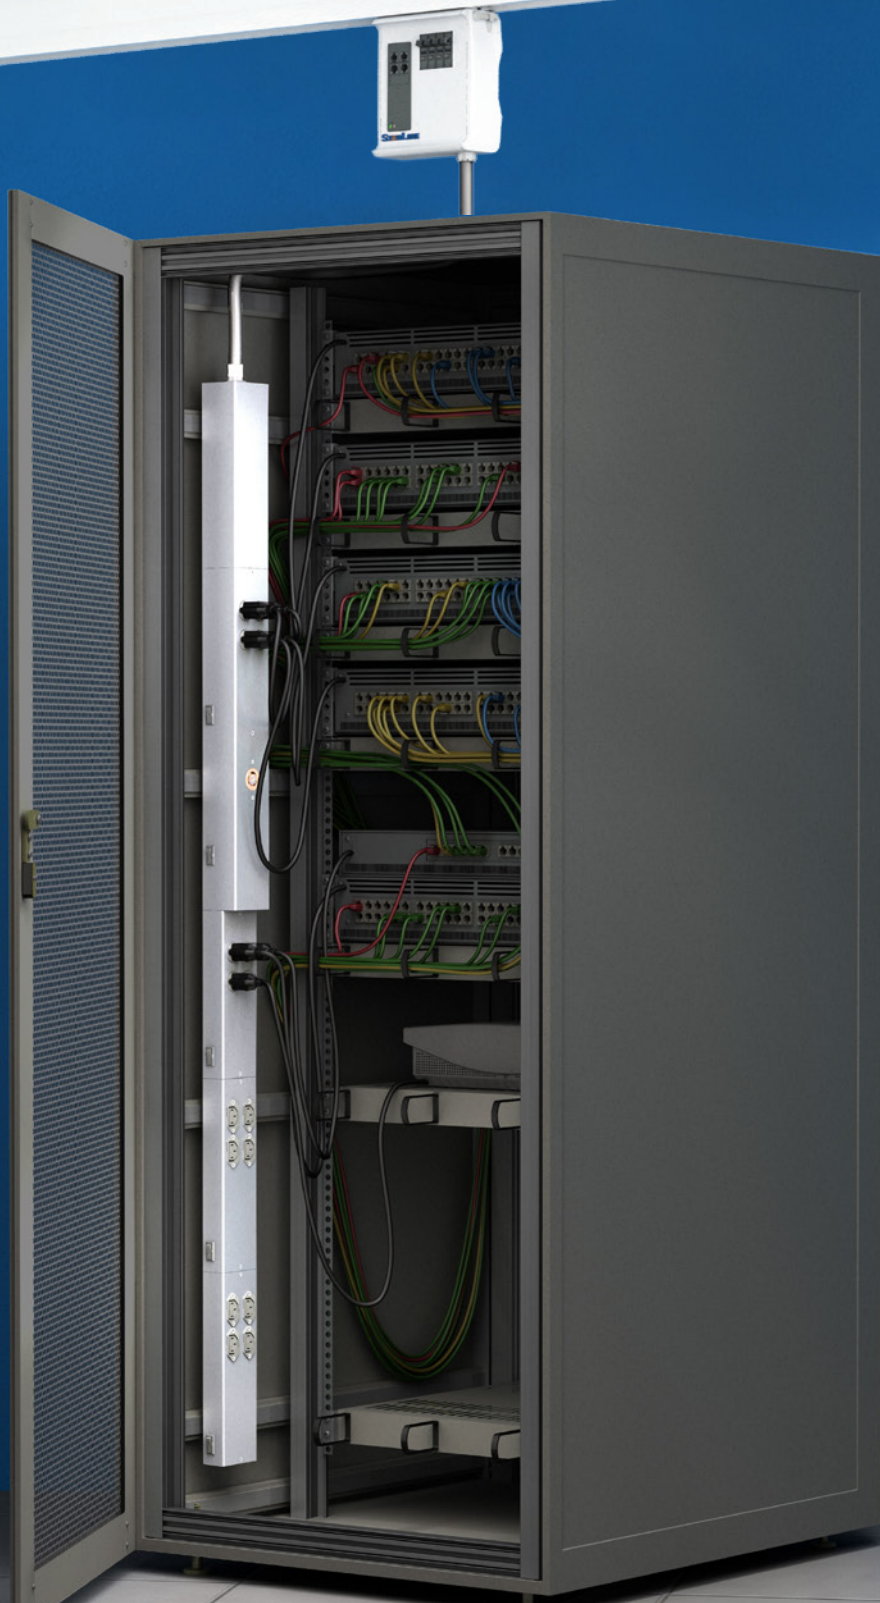

E+I ENGINEERING GROUP

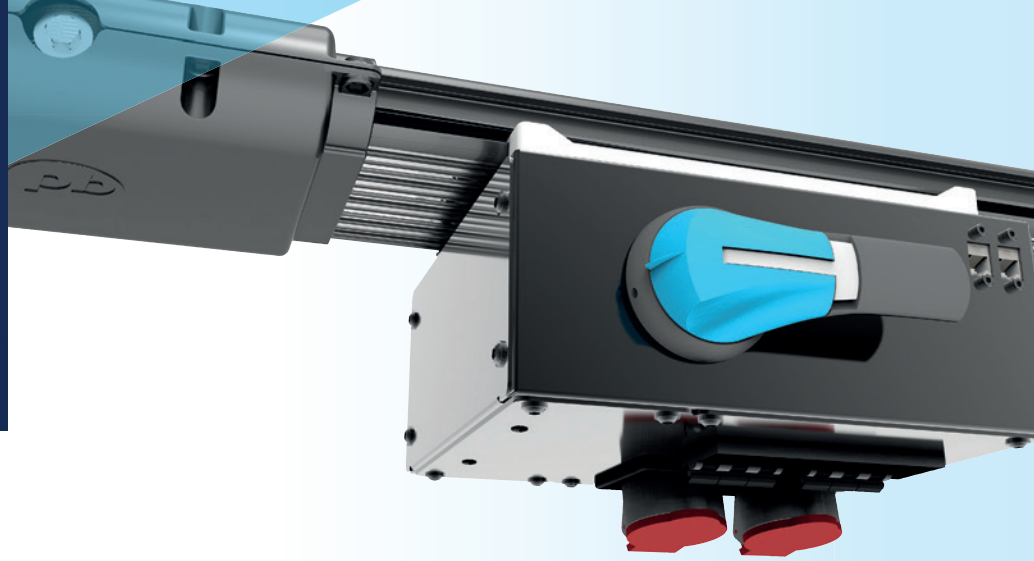

DELIVERING POWER SAFELY AND EFFICIENTLY IN MISSION CRITICAL ENVIRONMENTS

E+I Engineering design, build and test a full range of power distribution equipment in our facility in South Carolina.

- **iMPB**
  The safest and most flexible open channel busway system on the market. **iMPB is the only UL listed busway that does not need to be de-energized to add or remove circuits**
- **Low Voltage Power Control Systems**
  Custom built UL 891 and UL 1558 switchgear designed to improve efficiency, enhance safety and save space
- **High Powerbar**
  A power solution for transformer to main panel board connections, available in ratings up to 6600A
- **Powerboard Series**
  High efficiency PDU and RPP systems providing the highest level of protection for your equipment

To learn more about our innovative and flexible technology please contact us at info@e-i-eng.com or +1 864 375 1757

South Carolina, USA | Donegal, Ireland | Ras al-Khaimah, UAE

WWW.E-I-ENG.COM

pb PowerBar

Fueling the future

Sebastian Moss looks into Microsoft's multi-year effort to use gas-powered fuel cells at the rack level

Sebastian Moss
Reporter

When deciding on a location for a data center, ensuring that the site has a reliable connection to the electric grid is normally one of the first things any company has to think about.

Microsoft believes it might have found another way. For the past five years, the company has been experimenting with a different approach - hydrogen or methane gas-powered fuel cells at the rack level. This offers advantages in terms of location, reliability and energy efficiency, but brings with it a whole new set of challenges.

After trials at Wyoming's Dry Creek Water Reclamation Facility and the National Fuel Cell Research Center at the University of California Irvine, last year saw Microsoft take its biggest step yet towards commercializing this idea.

In September, the company teamed up with McKinstry and Cummins to launch what it called "the world's first gas data center," the design of which went on to win the DCD Mission Critical Innovation Award, presented at the end of 2017.

"The Advanced Energy Lab is a 20-rack data center pilot located in Seattle," Microsoft research manager Sean James said.

"What makes this project so disruptive is how radically it simplifies the process of powering servers and how this could almost double the energy efficiency of data centers - all while reducing costs and improving reliability."

The project design, known as 'The Stark and Simple Data Center,' puts fuel cells at the server rack level. A fuel cell is composed of an anode, a cathode and an electrolyte layer between them.

In the case of hydrogen-fueled proton-exchange membrane fuel cells (PEMFCs), hydrogen molecules split at the anode into protons and electrons, activated by a catalyst. The protons can pass through the membrane while the electrons cannot, causing them to travel across an external circuit, therefore producing energy.

Different fuel cells use different electrolytes and charge-transferring ions, but the two with the most promise for the data center appear to be PEMFCs and solid oxide fuel cells (SOFCs), with the latter using natural gas (methane) directly.

SOFCs have a higher efficiency, larger capacity, and are suitable for continuous power generation, but run at a higher temperature, produce carbon dioxide, and are less responsive to variable load.

PEMFCs remain the cleanest fuel cell variant, producing only waste water, but hydrogen distribution and storage has yet to scale. That could change in the future, however (see *Dawn of the hydrogen era* on p23 for more).

SOFCs present a quick and viable alternative to the electric grid and can still provide some environmental advantages over combustion - Microsoft estimates carbon dioxide emissions could be reduced by up to 49 percent, nitrogen oxide by 91 percent, carbon monoxide by 68 percent and volatile organic compounds by 93 percent.

Microsoft is experimenting with both approaches - in the paper *Fuel Cell Powered Data Centers: In Rack DC Generation*, the company explained that it has run tests using both a Hydrogenics HyPM 10kW PEMFC and a SolidPower Engen 2500 2.5kW SOFC.

The Hydrogenics system is designed with data centers and telecoms facilities in mind, whereas similar products do not exist with SOFCs just yet. The SolidPower system targets the residential and commercial buildings sector, and was designed for the co-generation of heat and power.

The paper states: "It is expected that SOFC systems designed for the in-rack power generation in data centers will produce more electricity and less heat and that such will be cooled with excess air and air ventilation systems already present in the data center."

Key to the method's energy efficiency improvement is the decrease in transmission losses. Energy is lost as it travels down the wire and, in a standard data center, these losses are found in the power plant, transmission, substation, transformation, and AC to DC conversion at the server. ▶

Microsoft's Christian Belady and Sean James

With fuel cells spread across the server racks, a method known as a Distributed Fuel Cell (DFC), the loss is limited just to the power conversion in the cell itself.

Microsoft claims that less than 35 percent of energy supplied by a power plant is actually delivered to the data center due to generation, transmission and distribution losses. When energy consumption associated with cooling, lighting and energy storage on the site is taken into account, only 17.5 percent of the energy supplied by a power plant reaches the servers.

In a research paper, *Fuel Cells for Data Centers: Power Generation Inches From the Server*, Microsoft found that when using its gas fuel cell approach, as much as 29.5 percent of generated power was spent at the server, assuming that the cooling load is proportional to the electric load for both traditional and fuel cell powered data centers. Further gains may be made if one could reuse the fuel cell waste heat for cooling processes.

Efficiency is not the only thing that makes a gas-powered data center attractive. Reliability could also be a major draw, with gas grid infrastructure usually buried underground and safe from severe weather. Gas grid reliability in the US is more than 99.999 percent, compared to the electric grid's 99.9 percent or less.

This means that data center operators can consider eliminating diesel generators and batteries used for power backup, as well as getting rid of the power distribution system in the data center, which can account for over a quarter of the capital cost of a data center.

Challenges remain, however, with the technology still in the research stages. Of course, whilst operators no longer have to worry about a reliable connection to the electric grid, such concerns are replaced by a need to locate near a robust gas supply, if the SOFC approach is taken.

But perhaps the biggest issue is that of rapidly varying loads. Servers can drastically change their energy requirements at a moment's notice, which the electric grid can easily accommodate.

But, with fuel cells, there is "a period of time where the fuel cells deliver an insufficient amount of voltage or power, resulting in server damage or shut down, which may lead to data center unavailability," Microsoft explained in the research paper *SizeCap: Efficiently Handling Power Surges in Fuel Cell Powered Data Centers*. "We call this phenomenon a power shortfall."

There are two approaches to dealing with such shortfalls. One is to perform power capping, which throttles server execution, therefore limiting server performance and causing both the fuel cells and data center to be underutilized. The other approach is to use energy storage devices (ESDs) like batteries and supercapacitors to cover the shortfall during load surges.

Provisioning for the worst-case scenario comes at a high financial cost, however. With this in mind, Microsoft has proposed SizeCap, which "coordinates ESD sizing with power capping to enable a cost-effective solution to power shortfalls in data centers." The hope is to find a middle ground between the two approaches.

Further work on SizeCap and the matter of fuel cells as a whole is necessary, and the cost of the cells remains prohibitively high. Microsoft is hopeful, however, that the method will soon see more widespread adoption. If that happens, concerns over data centers' impact on the electric grid may become a thing of the past.

Dawn of the hydrogen era

While proponents of solid oxide fuel cells (SOFCs) will point to the vast network of existing gas infrastructure that spans the globe as one of its greatest benefits, there are those that are adamant the future lies in hydrogen-powered proton-exchange membrane fuel cells (PEMFCs).

"I actually see SOFCs as more of a niche market, and it's fairly expensive," Dr Steve Hammond, director of the Computational Science Center at the US Department of Energy's National Renewable Energy Laboratory (NREL), told *DCD*.

"They are best suited for baseload power and variable load or load following is not their strength."

PEMFCs, meanwhile, have experienced significant advances "in terms of the efficiency and cost which is being driven by the automotive industry," Hammond said. "I think we can leverage that for consideration in the data center space."

While a separate battle rages between hydrogen and lithium-ion battery powered cars, manufacturers have invested heavily in the technology, helping bring the price of fuel cells down considerably.

Already, automotive firms are taking steps to use this technology for the data center. Late last year, Daimler revealed it had partnered with Hewlett Packard Enterprise (HPE), as well as Mercedes-Benz and NREL, to develop prototype continuous power solutions and stationary power systems that rely on hydrogen.

"The maturity of automotive fuel cell systems is unquestioned today. They are ready for everyday use and constitute a viable option for the transportation sector," Prof. Dr Christian Mohrdieck, fuel cell director at Daimler AG and CEO of Daimler subsidiary NuCellSys, said at the time.

"However, the opportunities for hydrogen beyond the mobility sector - energy, industrial and residential sectors - are versatile and require the development of new strategies. Economies of scale and therefore modularization are important challenges."

NREL has been testing a prototype of Daimler's system since November 2017, installing it at its Energy Systems Integration Facility (ESIF) data center. Purported to be one of the most efficient sites in the world, it had an annualized average power usage effectiveness (PUE) rating of 1.04 in 2016.

"It's still in the basic research phase, we're focused on 'can we do it safely, can we do it at all, does it make sense, is there a path to economic viability and where is further R&D required?'," Hammond said.

"We've had a prototype, I believe Daimler plans to provide us one of their actual production versions that we'll be able to operate by mid-year."

That will represent an early step, with the technology still essentially repurposed automotive gear. "The things we're doing are retrofitting automotive fuel cells, there's currently no 500-100KW purpose-built data center specific fuel cells to date," Hammond said.

"I think if it's going to have an impact on the cloud space, it's got to be done at row-scale integration. We have worked with Microsoft and gone through some techno-economic analysis of what a 50MW data center would look like if it were off-grid. They want 100 hours of backup. We went through the analysis at two different locations in the US, and how much wind and solar would be needed to support, and how you might store five gigawatt hour equivalents of hydrogen," he added.

"They're interested and they're looking, and at some point it will make economic sense, it's not immediate."

Economic sense is not the only reason for cloud companies to be attracted to hydrogen fuel cells - there's also the possibility of significant emission reductions. The extent of these reductions depends, however, on how the hydrogen is harvested.

The production of hydrogen always requires more energy to be put in than can be retrieved from the gas, in a process that relies on fossil fuels - some 95 percent of all hydrogen fuel is created using natural gas.

The steam-methane reforming process, currently handled by large-scale reformer processing devices, emits carbon dioxide and carbon monoxide as byproducts.

To be truly green, the hydrogen must be produced with electrolysis, where electricity is run through water to separate the hydrogen and oxygen atoms - and that electricity must come from renewable sources.

"For us to demonstrate that we can do this at all, we have solar on campus," Hammond said.

"Through purchase power agreements, that feeds our electrolyzers, and we have existing hydrogen storage infrastructure. So we wanted to show that it was possible at all and possible to do safely, with appropriate safety mechanisms and failure mechanisms to show that - my biggest fear is that I didn't want to be that other hydrogen incident people refer to."

The Hindenburg disaster that killed 36 people in a devastating inferno in 1937 did not just mark the end of the airship era, it also taught an entire generation to distrust hydrogen. "People ask about the volatility of hydrogen, but they're perfectly comfortable driving around in an automobile with 20 gallons of gasoline," Hammond said.

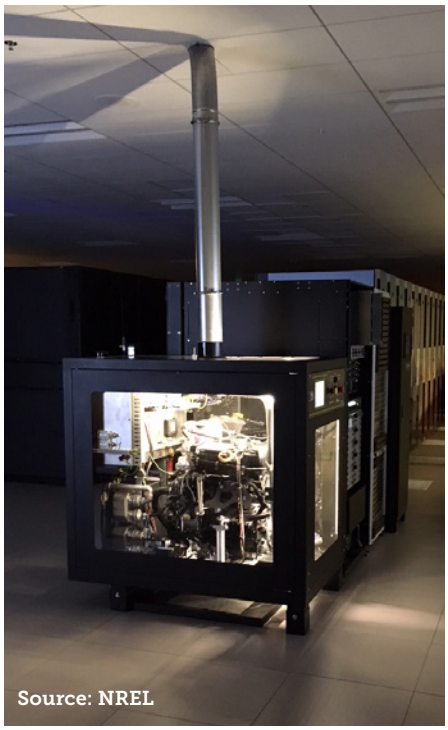

Source: NREL

1838
The invention of the first fuel cell

"At NREL we've worked a lot on hydrogen sensors and technology for fueling stations. We work with hydrogen all the time, perfectly safely."

But, for data center technicians to end up doing the same will still require many more years of research from those interested in radically rethinking the architecture of a data center, and how power is produced and stored.

For the past three years, Hammond has been working on the concept without direct funding. "I see this as impacting the market in five to ten years," he said.

For more on the subject, Dr Steven Hammond will be presenting 'Fuel cell innovation in the data center - can it replace the grid?' at DCD Energy Smart in Stockholm on March 13.

To register, or see the full conference agenda, go to: bit.ly/DCDEnergySmartInfo

Why Pkaza?

Contingency Recruiting with Guarantee
Largest listing of data center facilities jobs in the USA
Core focus is recruiting for the data center facilities industry
Highly Networked with many of the best companies and candidates

We Recruit for both End Users and Vendors in the Data Center Market

RECRUITING EXPERTISE

MEP Engineering & Design	Electrical / Mechanical / Controls / FPE / Cable RCDD / AutoCAD / PMs
Construction	MEP Supers / Construction PMs / Project Executive / Owner's Reps
Commissioning / CXA	Electrical / Mechanical Commissioning PMs / Agents
Sales	Solutions / Product / Service / Consulting Sales Professionals
Operations / Maintenance (O/M)	Directors / Chief Engineers / Facility Managers / Engineers / Techs
Field Service	Electrical / UPS / HVAC - Technicians / Service Managers
Executives	Director, VP, President, and C-Level Execs

ATTENTION: HIRING MANAGERS!

Are you settling when hiring for talent?
Use Pkaza for your niche data center recruiting needs!
You don't pay a dime unless you hire one of our candidates!

Pay It Forward! Not looking but know someone who is? We pay for referral$!

Contact us at:

peter@pkaza.com
973-895-5200
www.pkaza.com

Manufacturing moves to the Industrial Edge

Digital infrastructure is changing the way factories operate - with far-reaching consequences for all of us, say *Sebastian Moss* and *Peter Judge*

Peter Judge
Global Editor

Sebastian Moss
Reporter

Factories are becoming smarter, more automated, more flexible and more efficient using new technologies. It's often called Industry 4.0, but there is more to this than a glib buzz-phrase. Data center technologies and other tools from digital infrastructure are being transformed and put to work in all vertical markets. We'll look at other sectors in future issues, but first lets see how manufacturing is being transformed.

Factories produce and use vast quantities of data, including specifications for the product, inventories of components, designs... and changes and variations to those designs. For best results, that data should be located and handled right there, near the factory itself.

Factory floors are managed with mission critical IT, housed in a specialized data center, most likely a containerized facility, inside the factory or just outside it - and connected to the cloud. Data center professionals will recognize this as an 'edge' resource. The data has to be stored and handled close to the the staff and robots of the factory, ensuring very low latency.

Factories are being transformed by digitization, so the micro-data centers they use are often being built with modular growth in mind.

Factory IT also has to be rugged, and immune to the dangers of the manufacturing environment, which include dust, pollution, vibration and electromagnetic pulses. Despite being distributed through the factory, it must be secure. And of course it must be highly reliable with redundant power and climate control just like one would expect of a mission critical facility.

Industry 4.0 is the fourth industrial revolution, turning factories into smart factories, filled with sensors that collect data to intelligently optimize operations, driving more and more automation.

Instead of simple, individual machines, the aim here is to create Cyber-Physical Systems - networks of connected machines that collect massive amounts of data ▶

Merging IT and OT

As well as eliminating a long-standing and unproductive turf war, the merger of IT and OT (operational technology) will have multiple benefits: it will allow IT managers to help mitigate security risks from networked mechanical equipment - which are emerging as a serious concern. At the same time, OT managers' expertise in physical machines will extend the ability of IT.

In manufacturing, the IT customer has historically been in a stronger position than the OT customer. With ERP software and other industry offerings seen as mature, switching products may be costly, but not impossibly so.

In OT, systems are often unique, task specific and proprietary, making it incredibly difficult to have market choice and price competitiveness. This may change with IT/OT convergence, with industry leaders from both sides coming together, and blurring the lines between the two groups.

through sensors, and then act upon that data in ways that are only possible due to this new level of insight and perception.

The factory processes can be integrated with the supply chain and even customer relations management (CRM) allowing vast increases in flexibility and customization.

Germany coined the term 'Industrie 4.0," and the world's largest manufacturer, China, is eager to be at the fore. In 2015, an Accenture report said that Industry 4.0 could add $14.2 trillion to the world economy over the next 15 years.

Robots are a big part of this: the Changing Precision Technology Company factory in Dongguan City, China, cut its workers from 650 to just 60, and its defect rate from 25 percent to 5 percent with the use of robots.

But networking is the real key. Industry 4.0 requires an Industrial Internet, and it also demands that information technology (IT) finally merge with operational technology (OT) - the branch of the organization handling mechanical devices and physical infrastructure.

IT and OT have managed to maintain separate domains in most organizations, but manufacturing relies so heavily on mechanical systems that Industry 4.0 cannot emerge without the two worlds finally coming to terms with each other.

As those in the industrial sector will know, manufacturing plants often contain legacy equipment, which might not necessarily possess the required digital connections.

These devices need different sensing strategies - some are non-intrusive, for example using an infrared camera to learn the temperature of a furnace. Other cases may require ad-hoc development, involving a piece of equipment close to the machine that collects data. The approach depends on the specific challenge and the cost.

It's clear that digital infrastructure is bringing a whole new approach to manufacturing while, at the same time, the industrial vertical is demanding new capabilities from sensors, edge data centers and micro-facilities.

For an extended report on Industry 4.0 visit: bit.ly/DCDManufacturing

Four Industrial Revolutions

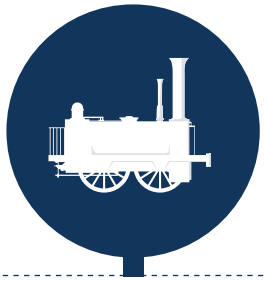

1.0 Steam and water

In the 18th and 19th centuries, factories emerged, using steam and water power, creating the first wave of automation

2.0 Electricity

At the end of the 19th century, electricity arrived as a more flexible power source

3.0 Computers

In the late 20th century, numerically-controlled machines such as lathes, and robot arms began working with data from computer-aided-design files

4.0 Full digitization

Beginning this century, all stages of the production process are automated and connected, factories are instrumented with sensors,and manufacturing becomes more flexible and efficient

Digital pathways

The data lifecycle begins on the factory floor, with Industrial Internet of Things sensors capturing readings from machines involved with the production process. Real time data from distributed control systems, manufacturing execution systems (MES), asset management systems (AMS), etc. is brought together to build digital twins of the machines on the factory floor.

Physical assets are modeled in data, creating a "digital twin" which can be fed with external data, allowing an understanding of what's happening inside the machine without stopping it.

Creating digital twins of factory machines and the supply chain itself can allow companies to keep a complete digital footprint of their products from conception to delivery. This allows for rapid process changes to be tested virtually, giving businesses new insights into how to improve their products, increase their yields and better serve their customers.

Sometimes serious issues can be spotted in this data allowing preventive action to be taken.

With this in mind, data is often stored at the edge, near the factory and the machines themselves. With latency minimized, errors can be avoided before humans are even aware of any issues, allowing for predictive maintenance, which is vastly more cost-effective than dealing with breakdowns after they have happened.

With edge data centers more powerful than ever, they can run their own private cloud that allows for the necessary processing (and protection). Once the perishable data is filtered out, some information can still be sent to the outside cloud and added to data coming from other factories and other elements of the supply chain.

Effectively managing the data journey is about mixing the real time electronics of the machines on the factory floor with the immediate response of edge facilities, and the big data insights that require massive computing power, which can be easily obtained in the cloud.


DCD > Enterprise | New York
16th Annual
Global Content Partner
Uptime Institute
Headline Sponsor
Life Is On Schneider Electric
Focus Day on EDGE | April 30
Free qualified end-user passes available
THE DATA CENTER & CLOUD INFRASTRUCTURE TRANSFORMATION SUMMIT
May 1-2 2018 // New York Marriott Marquis
Limited sponsor and exhibitor opportunities - contact chris.hugall@datacenterdynamics.com
For more information visit www.DCD.events
@DCDConverged #DCDEnterprise
Datacenter Dynamics
DCD Global Discussions

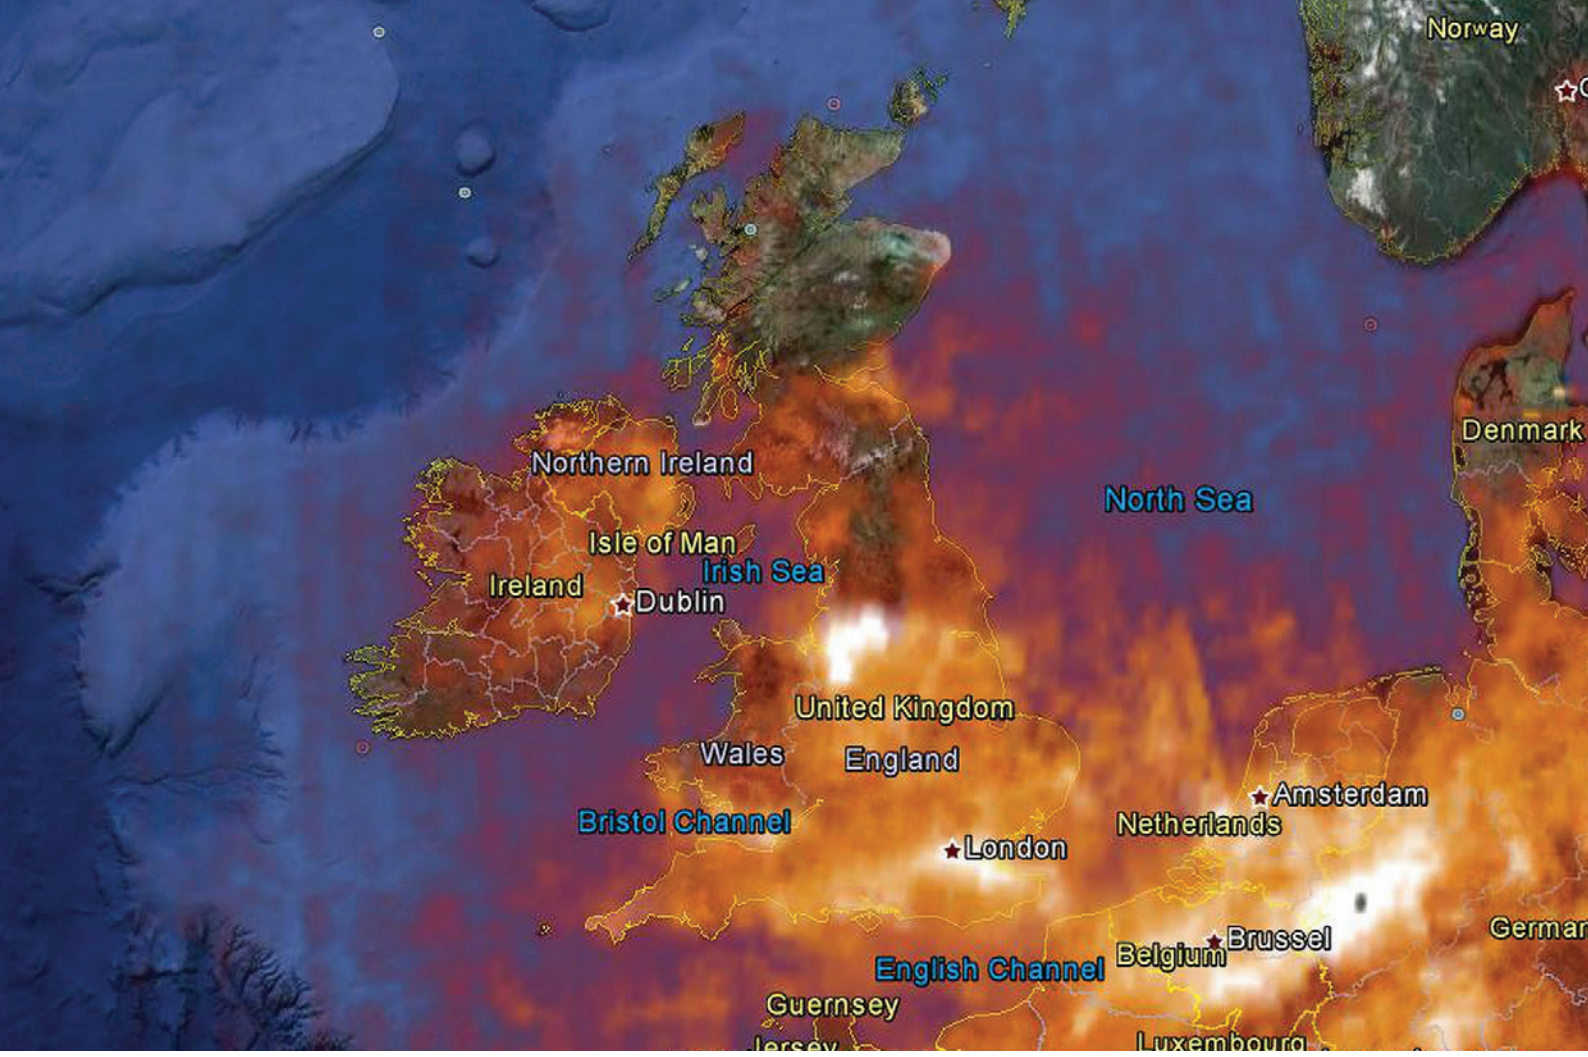

Clean air cuts costs

Equipment in data centers has become more sensitive to airborne pollution, thanks to environmental regulations and the use of outside air for cooling, say *Paul Finch* and *Chris Muller*

Paul Finch
KAO Data

Chris Muller
Purafil, Filtration Group

The IT equipment lifecycle may be as short as three years, but it is still worth maximizing the life of servers, storage and other electronic equipment within the data center. Unfortunately, airborne contamination can cause hardware failure in servers, hard drives, UPS systems and other sensitive electronic equipment.

Contaminants can be introduced during infrastructure upgrades, or by simply opening doors while entering or exiting a white space.

These contaminants affect electronic equipment, corroding components and reducing capabilities to a point of failure, which can result in costly but often avoidable data center outages.

'Free-cooling' with outside air can exacerbate the contamination issues due to the relatively large volumes of air required to maintain the environmental conditions within the data center.

Even when chilled water, condenser water and indirect outside air systems are used, contaminants can still be present, as building regulations and codes dictate that nominal volumes of outside air must always be provided.

Because of the latency and bandwidth requirements of the business, some data centers must be located within highly populated areas with heightened pollution levels, or in industrial parks where neighboring businesses may generate pollutants.

Even data centers located in areas without air quality concerns must plan to ensure they can maintain an environment during the lifecycle of the facility that provides protection of sensitive electronic equipment.

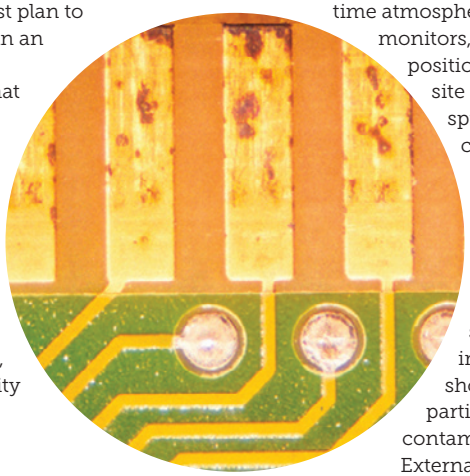

Sites where the server environment is maintained at especially high temperatures must be constantly monitored for changes in humidity, because variable humidity increases the risk of corrosion.

Some environmental regulations have had unintended consequences, such as the European Union's Restrictions on Hazardous Substances (RoHS), which eliminated the use of lead in manufacturing of electronic components. RoHS requires the use of lead-free alloys in solders and PCB surface finishes which makes them far less resilient to deterioration in contaminated environments.

RoHS-compliant IT and datacom equipment is at risk in locations with poor ambient air quality, as products using an immersion silver (ImAg) surface finish will corrode in environments with a high sulfur content. Data centers in urban areas have reported server, hard disk and UPS system failures caused by sulfur corrosion.

Manufacturers have set specific working parameters and contaminant levels for their products to remain within warranties, and ASHRAE's Technical Committee 9.9 (TC9.9) has issued guidelines on gaseous and particulate contamination limits for data centers, which were incorporated in ISA Standard 71.04-2013.

However, data on contamination levels can be unreliable, often depending on such highly-variable factors as pollutants from motor vehicle exhaust.

For this reason, operators must monitor and control environmental contamination within their data centers to understand the real threat to their equipment.

This should be part of all data center management and planning practices, and include:

- Assessment of the outdoor air and indoor environment
- Development of a contamination control strategy
- Real-time environmental monitoring

Airborne contaminants in a data center can be measured using 'reactivity monitoring' which measures copper and silver corrosion rates using specially prepared sensors correlated to ISA severity levels. Real-time atmospheric corrosion monitors, placed at specific positions around the site can measure the spread of airborne contamination.

External and internal filtration systems should be installed in existing or new air conditioning systems without impairing airflow, and should remove both particulate and gaseous contaminants.

External filters must deal with PM10 (particulates between 10µm and 2.5µm), PM2.5 (particulates smaller than 2.5µm) and ozone, and incorporate chemical filtration for local contaminants such as motor vehicle exhaust and industrial activities.

Internal filters must also provide for chemical filtration. A large contact area is required to increase the efficiency of the filter and a parallel structure offers unobstructed airflow to ensure a low pressure drop.

Monitoring systems should use wireless communication to gather real-time site-wide data, including temperature, relative humidity and room pressure, along with confirmation of the copper and silver ISA severity levels. If they are above the IT and datacom equipment manufacturers' warranty requirements, steps should be taken to reduce and control the amount of corrosive gases in the data center environment.

Contamination measurement and control is a necessity due to the tightened operating specifications for today's datacom and IT equipment, and an increase in sensitivity of electronic devices caused by RoHS and similar regulations.

Outdoor air used for ventilation and pressurization must be cleaned as manufacturers' warranties demand a clean environment.

There is a cost to this, but the effects of airborne contamination will be far more detrimental. After all, an ounce of prevention is worth a pound of cure. ●

Paul Finch is chief operating officer at KAO Data; Chris Muller is technical director at Purafil, Filtration Group

This is an abridged version of an article on DCD's website bit.ly/DCDCleanAir

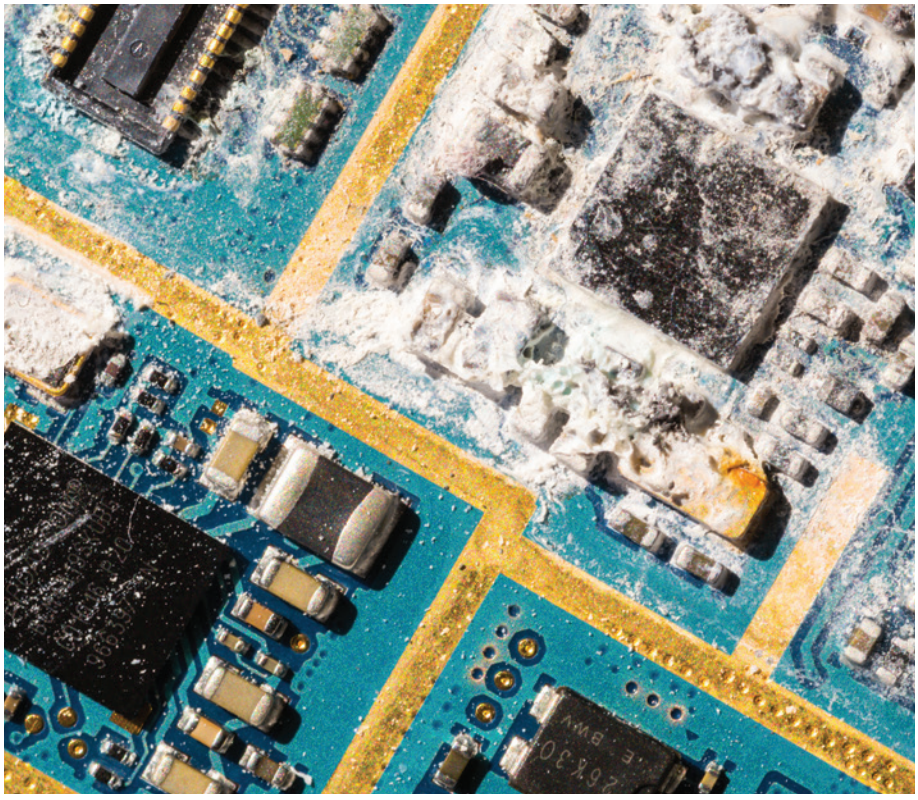

After the Storm

Sebastian Moss investigates the supercomputer that saved Cray, shot down a satellite, and gave birth to a generation of HPC systems

Sebastian Moss
Reporter

The history of high performance computing (HPC) can be told through the stories of a few key supercomputers whose impact reverberated through an industry marked by the rise and fall of companies racing to build the most powerful systems in the world.

Red Storm is one such system. When the supercomputer was decommissioned in the summer of 2012, the president and CEO of its creator, Cray Inc, summed up what made this particular machine so important.

"Without Red Storm I wouldn't be here in front of you today," Peter Ungaro said. "Virtually everything we do at Cray - each of our three business units - comes from Red Storm. It spawned a company around it, a historic company struggling as to where we would go next. Literally, this program saved Cray."

To understand what Ungaro meant, we must go back to February 1996, when the American supercomputing manufacturer was acquired by Silicon Graphics (SGI) for $740 million.

Cray was working on the T3E, its second-generation massively parallel supercomputer architecture (after the T3D) and "SGI basically took most of the folks that were working on the T3E project and redirected them to extending the microprocessor-based line that SGI was building," Cray's chief technology officer, Steve Scott, told *DCD*.

Four years later, "SGI spun Cray off and they kept all of the parts of the company that were working on massively parallel processors using microprocessors," Scott said. "The rest of us became Cray Inc, which is the company that we are today."

The new entity, which was "relatively small at the time, with revenues of under $200m," was left working on the X1 and X2 systems - commercially limited products propped up with funding from the NSA - and faced the very real possibility of extinction.

"So that was the environment that we were in in 2001 when Sandia National Laboratories first started talking to us," Scott said.

But the relationship between the US Department of Energy's Sandia Laboratories and Cray got off to an inauspicious start. "The first meeting between Sandia and the bidders - there were two - was on September 11th, 2001," John Noe, department manager at Sandia, told *DCD*.

The laboratory, one of three National Nuclear Security Administration research and development facilities, was put on high alert. "That was an interesting day, because we'd started the meeting on the base, and eventually we had to be taken off the base," Noe said.

"And neither one of the bids was what we were after in the end, it didn't look too promising," he added. Cray had tried to pitch its X1 and X2 line, but Sandia wasn't interested, telling them that linear processors were not the future.

With airports closed, "a bunch of us ended up driving home in a rented van across the country from Sandia to Wisconsin after that," Scott said.

The fate of Cray, a company whose founder Seymour Cray was known as the father of supercomputing, hung in the balance.

Over the course of a year, and after a lot of meetings, Sandia and Cray slowly grew closer, ultimately agreeing to develop what would become Red Storm - a system more akin to what the old Cray had achieved with the T3D and T3E and what Sandia and Intel had done with their ASCI Red System, the first computer built under the Accelerated Strategic Computing Initiative.

284.16 teraflops

The peak performance of Red Storm from 2008

But, while Red Storm was planned to be a massively parallel processing machine in the same vein as ASCI Red, that system had left Sandia worried that it was witnessing dangerous consolidation in the HPC marketplace that would reduce the amount of competition, and leave any supplier vulnerable to the whims of large customers.

"Intel's supercomputer system division had basically gone out of business as they delivered ASCI Red to us and it was the last and only of its kind," Noe said. "The lesson learned from that was that we wanted to make a commercial platform available, based on our technology, and so as part of the contract we made sure that Cray was going to sell it.

"We wanted to continue creating choices in the marketplace, having more than just one vendor option to go to, in the future, was part of our thinking."

Dave Martinez, project lead for infrastructure computing services at Sandia, concurred: "At the time there was IBM, and I think Red Storm put the competitive nature back into the HPC program and it modernized the way partnerships are formed and how projects are accomplished.

"It was a milestone," he told *DCD*.

With Cray in a tough place financially, this meant that Sandia had to rewrite its approach to contract negotiations. "The money they were getting from Red Storm essentially was their only cash flow," Noe said.

"We had to work very closely with them to make sure they were still viable and able to pay their people and buy their parts and so forth, whilst making sure we still received some value in the case something went wrong."

Instead of the usual small payment upfront and large payment at the end, Sandia "assigned intermediate deliverables to pay Cray, which kept them afloat and gave us some confidence that we weren't sinking a whole lot of money into something that we wouldn't get value back for," Noe said.

With the contract signed, both teams got to work: "The whole system was designed, really, from architecture to delivery, in just under 20-24 months," Scott said. "And this is really quite remarkable because it normally takes a lot longer than that to design a system."

As they planned out the project, Sandia had to simultaneously prepare the shell to house the upcoming system. "We had no idea, when we built the data center, what kind of equipment we were going to place in there, so we basically built the shell for a relatively low price," Martinez said.

"Not knowing what the future might bring, we decided to build it with larger pipes, thinking about liquid cooling in the future, and possible low airflow pressurization types."

But Red Storm ended up relying on air cooling. "Back in the day we just had centrifugal fans, and they basically throw air like 60 feet - there was like 80 feet to the middle of the system."

Martinez found inspiration at a local power company that was experimenting with plug fan technology. "It was a first one of its kind," he said. Working in partnership with HVAC company Johnson Controls, Martinez flew to Ireland to test out air conditioning firm EDPAC's fan units and came up with a concept of pressurizing the floor, working all of the CRAC units in unison, circling around the Red Storm computer.

The launch of the SM-3 missile that intercepted USA-193, programmed by Red Storm

The motor on the unit was not designed with that in mind, however, so Martinez flew to Milwaukee to use ABB's latest drive, which was still in the test phase. After some tinkering, telling EDPAC to change some of its "control schemes to cold site control instead of return air," and instructing them to adopt ABB's drives, Sandia found success.

"That's one of the first big data centers to be cooled by plug fan cooling," Martinez said. "Now, basically 95 percent of big data centers, if they have CRAC units, are cooled by using plug fans. The centrifugal fan has been replaced because it's not as accurate at cooling as a plug fan."

This concept allowed Sandia to cool "the largest powered machine we'd ever seen," at about 40kW per rack, "quite comfortably."

With a shell built, Sandia and Cray started installing the system and linking the components together. "It was a 3D Torus interconnect with a custom router chip called SeaStar that Cray designed. That network was using a little over three gigabits per second signaling rate," Scott told *DCD*. "So it was competitive with the InfiniBand networks of the day." ▶

Inside Red Storm

Red Storm used AMD Opteron processors, with an initial installation of 140 cabinets, taking up 280 square meters (3,000 sq ft) of floor space, operating on "a microkernel that was largely developed by Sandia with Cray involvement called Catamount."

With a fast turnaround of less than two years, the system "had a very... let's call it 'fresh' software stack," Scott said. "And so it was not the most stable machine and we proceeded to spend the next two years hardening it into a much more production-worthy machine."

He added: "Eventually, for example, the microkernel evolved to using Linux and we developed what's known as Compute Node Linux which is essentially the line of scalable operating systems that we still ship today." This operating system, Scott said, allowed Cray to corner the lucrative HPC market in weather forecasting, which requires high operational reliability.

Red Storm also saw several hardware improvements. It launched with a theoretical peak of 41.47 teraflops in 2005, was upgraded to 124.42 teraflops in 2006 and to 284.16 teraflops in 2008, as AMD quickly iterated its Opteron line from single, to dual and finally to quad-core processors.

"Because of the scalability of the design, it was certainly one of the seminal events in our growth from teraflops to petaflops," Noe told *DCD*. "We effectively demonstrated that with Red Storm."

The adoption of Opteron chips also served as "a bit of a wake-up call to Intel," Scott said, "they were really not focused on meeting the needs of the HPC community during the early 2000s, while AMD moved to 64-bit addressing, integrated memory controllers and integrated network ports. All three of those things were very soon thereafter taken up by Intel."

As it silently set about reinvigorating competition in the chip and supercomputing space, Red Storm was also working on other things - secret things. The computer's classified work is described by the laboratory as solving "pressing national security problems in cyber defense, vulnerability assessments, informatics (network discovery), space systems threats and image processing," as well as simulations involving nuclear weapons research.

One classified project was made public, however - Operation Burnt Frost.

After the US lost control of a military reconnaissance satellite soon after its launch in late 2006, concerns were raised over whether it would accidentally spray a cloud of highly toxic hydrazine fuel over whichever country was unlucky enough to serve as USA-193's crash site.

It took until January 8, 2008 for the Burnt Frost program to be created, with a decision to shoot down the satellite taken soon after. The US Navy and the Aegis Ballistic Missile Defense then had just a few weeks to pull it off. They turned to Red Storm.

Sandia simulated thousands of variations of strikes against the bus-sized satellite traveling 17,000 miles per hour, 153 miles above Earth, before finally getting the approval from President George W. Bush to proceed with a takedown. Launched on 20 February, the strike was a total success - although it received

1976

When the first Cray supercomputer was installed at Los Alamos National Laboratory

criticism from China and Russia which claimed the whole exercise was a cover-up to test an anti-satellite weapon.

Lt. General Henry A. Obering III, the director of the Missile Defense Agency, said in a promotional video for the initiative: "Not only can we hit a bullet with a bullet, we can hit a spot on the bullet with a bullet."

Red Storm was also involved in unclassified projects, including climate modeling. "They wanted to be absolutely certain that there was no way for code running in the open partition to read any data or see any activity or code that was running in the classified position partition," Scott said.

"And so they asked us to build a couple of large sets of switches that would allow us to mechanically disconnect the interconnect."

With previous systems, such as Intel's Paragon, moving between classified and unclassified was a risk because they were not "designed to be switchable," Noe said.

"We had to do some operational things to be able to move it back and forth between different security realms - and, unfortunately, that involved disconnecting pieces and parts and turning it off, which is never a good thing to do with a big computer. So we would lose component parts every time we tried to switch it back and forth."

To address this, Sandia experimented with backplane disconnect units for ASCI Red, but took the idea a step further with Red Storm. The laboratory and Cray created a 'Red-Black Switch,' which allowed the system to fully run classified or unclassified material, or be split to run each simultaneously - with a physical 'air gap' ensuring security.

"The reconfiguration architecture that Red Storm embodied was quite remarkable," Noe added.

After a long life of both classified and unclassified simulations, the system was decommissioned in 2012. But its legacy is still visible in the industry.

As part of its efforts to ensure the HPC market stayed competitive, "the contract specifically called for Cray to sell a system derived from Red Storm as a commercial product," Noe told *DCD*.

And, indeed, Cray did - selling the system as the commercially successful XT3. That was followed by the XT4 and XT5, before being replaced with the XE system in the late 2000s and the Cray XC system in 2012, which is what the company ships today.

Developed with an initial budget believed to be around $72m, Red Storm's successors' sales have gone on to raise billions.

"Red Storm was the system that has allowed Cray to grow to several times the size that we were, and be successful over the years," Scott said. "It all started with this system."

As for Sandia, after Red Storm, it was left with a data center built with liquid cooling in mind and with three foot of raised floor space. "Since then we've placed several different systems in there. We're able to do liquid cooling, and air cooling in the same venue," Martinez said.

"We bought a couple of other air-cooled systems and we built some nice containments over them. We designed our own containments and we deployed them here. We started doing air-side containment before it was really out there."

Looking back at Red Storm, Martinez told *DCD*: "It was a great system. Now, we have other systems that we're working on that I can't say much about that equal Red Storm in impact at its time."

Dave Martinez and Cray's liquid cooled Sky Bridge supercomputer

Sandia's hot air balloon curtain

An air-side aside:

The first time Sandia National Laboratories tested air-side containment was on the Thunderbird supercomputer, a system built by Dell with 4,480 dual-Xeon servers, but its approach was far from conventional.

After installing the 60 teraflop HPC system in 2006, the laboratory soon realized that it had trouble cooling its new machine.

"We put the right amount of air to supply it," Sandia's Dave Martinez told *DCD*. "But the problem was, even with the plug fans, the air was going by so quickly that the bottom four to five nodes of each rack were overheating."

His solution was to get a "piece of plastic and put it over the top to take that air, hit the top, and let it roll down."

After a short test and some thermal imaging, it was clear that the concept worked, but a plastic sheet was far from a long-term solution. "The next thing to do was to find what type of material we should use that the fire marshal would allow us to put in there."

As luck would have it, "we are the balloon capital of the world and my then-manager had a hot air balloon.

"Using the hot air balloon material, we built it over the top of the racks, where it's just on a rail system, and could assemble it over two hours once it was all sewn on. So then we had our first containment system at Sandia."

Thunderbird was eventually replaced by a more powerful system, Red Sky, which increased rack density from 13kW to 32kW. That supercomputer relied on cooling doors from Sun Microsystems and Emerson Network Power (now Vertiv) to achieve a claimed power usage effectiveness (PUE) of 1.035.

The Sun rear cooling doors were 100 percent passive and relied on the inert refrigerant gas R134. But while it was an improvement over the Thunderbird, Martinez's hot air balloon quick fix - now over a decade old - still manages to find its way into American government HPC data centers.

"Our counterparts at Los Alamos borrowed our curtain about eight months ago as they needed to use a quick containment, and now the curtain is over at Lawrence Livermore and they're using it as a quick solution until they get something built.

"It stands up so quickly and is so effective, it has made its way around."

The many flavors of virtual

Dan Robinson
Correspondent

Virtualization can mean several different things, from bare metal to containers, says *Dan Robinson*

When Amazon announced it was adding bare metal instances to the portfolio of server options available to users of its Amazon Web Services (AWS) cloud platform, the move sparked much speculation. Would it have an impact on colocation and hosting providers, for example, as physical servers have traditionally been the domain of those companies?

Amazon unveiled its new Bare Metal service as a public preview at its re:Invent conference last November, enabling users to access an entire physical server rather than the virtual machine instances it mainly offers. Virtual machine instances typically offer just a subset of the capabilities of the physical machines they are running on.

So why should users bother with virtual machines at all? Because they are more flexible. Many cloud service providers operate a "pay-as-you-go" model, whereby users can provision a virtual machine when required, then halt or retire it when no longer needed, and pay only for the resources used when the machine is running.

Why are Amazon's bare metal services different from the physical servers which users can own and operate in colocation providers' facilities? Because cloud service providers take care of operating and managing the physical servers and all supporting infrastructure. Under the colocation model, the user organization usually has to purchase and manage their own equipment and rent space in the provider's data center to host it.

Amazon also claims that its Bare Metal instances have the same elasticity and scalability as other cloud instances, meaning that customers can provision them in minutes and scale them up or down as is possible with existing types of instances.

AWS is not the first cloud provider to add a bare metal capability to its portfolio – Rackspace, Oracle and IBM have offered bare metal servers for several years. However, Amazon's Bare Metal servers are notable for including dedicated hardware (a platform called Nitro) to offload network and storage handling in order to make as much as possible of the bare metal performance available for the user's application.

Customers have a choice between physical servers, bare metal servers and virtual machines, along with another option which is currently fashionable. Containers allow workloads to be deployed with a smaller subset of the functions of a physical server.

Virtualization on servers is a way of dividing up system resources so that multiple users or workloads can run independently on the same system without interfering with each other, and its roots go back to the first multi-user systems of the mainframe era. For example, IBM's VM (virtual machine) mainframe operating system was launched in 1972.

Virtual machines came to x86 servers came courtesy of VMware, when it unveiled its ESX Server and GSX Server products in 2001. On x86 servers, the main driver was workload consolidation. Before virtualization, corporate servers would often be running single workloads and operating at utilization rates as low as ten percent, meaning they were idle much of the time. Converting them into virtual machines meant that several loads could be operated independently side by side on a single physical server, reducing the overall number of servers an organization needed to have.

GSX Server allowed users to operate virtual machines on top of an existing operating system such as Windows, and is thus an example of a Type-2 hypervisor, whereas ESX Server ran on the bare metal and is thus an example of a Type-1 hypervisor. ESX Server gave way to ESXi, the compact, dedicated hypervisor that underpins VMware's vSphere platform today.

Type-2 hypervisors are less efficient, because they run on top of an existing operating system. For this reason, they are now largely restricted to client virtualization such as VMware Workstation, which lets a developer run one or more virtual machines on Windows or Linux desktops and laptops.

Other hypervisors commonly used for server virtualization include Xen, developed as an open source project in 2003. This was adopted by AWS to drive the virtual machine instances on its EC2 cloud service and is also the foundation for Citrix XenServer and XenDesktop.

A few years later in 2007, Linux gained its own hypervisor, in the shape of the Kernel-based Virtual Machine (KVM) project. This is so called because KVM is implemented as a kernel module that converts the Linux kernel into a bare metal hypervisor when loaded. Because it is effectively part of the kernel, KVM has grown to become the default option for many platforms that use virtualization, such as the OpenStack cloud framework, Apache CloudStack and most of the major Linux distributions.

The other major hypervisor in use today is Microsoft's Hyper-V. This has been built into every version of Windows Server since 2008, and is thus widely adopted since the vast majority of organizations use Windows servers. Many organizations that

Varieties of virtualization

- **Physical servers,** which run in private facilities or in colocation spaces.
- **Bare metal servers,** where cloud providers offer access to physical servers
- **Virtual machines,** a service provided which emulates a machine so multiple workloads run on the same server
- **Containers,** which offer a smaller subset of the functions of a server, so applications can run in a smaller and more portable environment

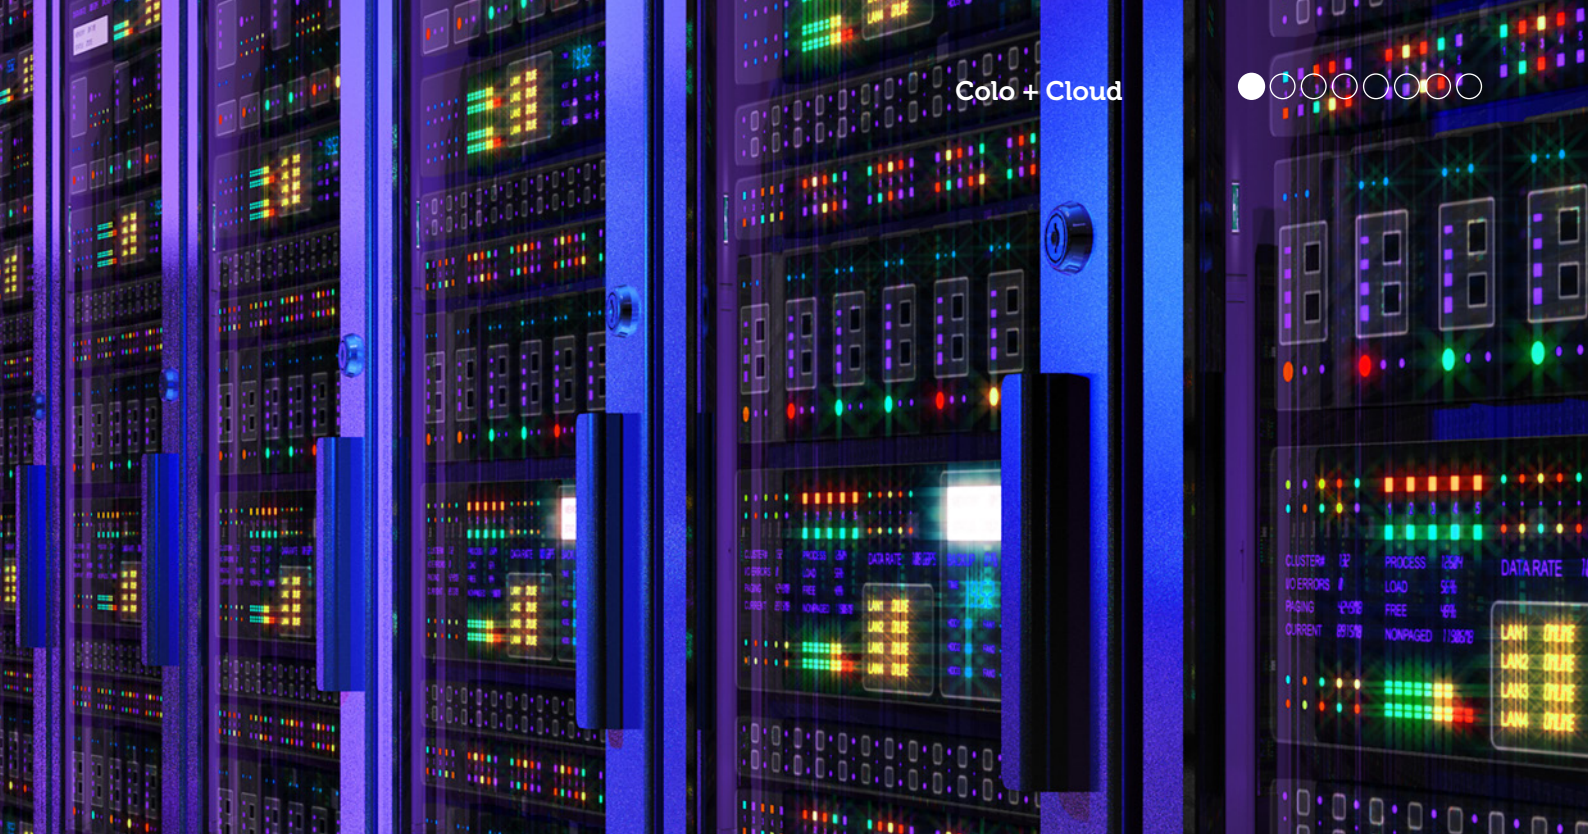

came late to virtualization will have been tempted to simply build upon what is already implemented in Windows.

All of these hypervisors differ slightly in the way that they operate, but the end result is much the same; they divide up the resources of the host server in order to create multiple virtual machines, each of which behaves as though it were a bare metal server in its own right. Users can even migrate virtual machine images from one hypervisor to a different one by using the right tools.

Each virtual machine must be provisioned with its own operating system and virtual disk before it can do any useful work, and may also need to be separately provisioned with the application that it is to run. In an enterprise environment, this will typically be handled by a management tool such as Microsoft's System Center Virtual Machine Manager (SCVMM), while in a cloud platform, this process is highly automated and driven by end users with a self-service provisioning tool or in response to some event.

Containers are an alternative way of dividing up system resources, but these operate at the level of the operating system. Instead of being an entire virtual machine, a container is essentially just an isolated environment within a server that contains application code and any supporting code libraries that application depends upon. As containers do not need to include an entire operating system, they can be quickly created and moved between servers, and it is possible to operate many more containers than virtual machines on any given server.

Containers have been around for many years, but the current explosion in container uptake is due to Docker, which in 2013 launched its namesake platform that enables developers to quickly deploy code inside containers. The Docker philosophy is to break down applications into smaller modules that can be deployed and updated separately, which melds well with current notions of agile development and a microservices architecture.

In the cloud, containers, virtual machines and bare metal servers are all just ways for service providers to sell units of compute power to customers. Each has its individual merits and use cases, and these are not mutually exclusive; a virtual machine could be used to host an array of containers that make up an application, for example.

Bare metal servers provide extra performance for demanding or specialist workloads, and allow users to deploy software that may be licensed only for non-virtualized environments, or to use operating systems that are not supported in the cloud provider's virtual machine catalog.

However, they are typically more work, as the user has to take on many tasks such as deploying and keeping the software up to date, which are managed by the cloud provider if you use a virtual machine.

The significance of AWS adding bare metal instances is that customers will be able to use these alongside the firm's virtual machine and container services, all from one console. It simply adds a new choice into the mix for customers to use as they see fit, and that is what the cloud is really all about. ●

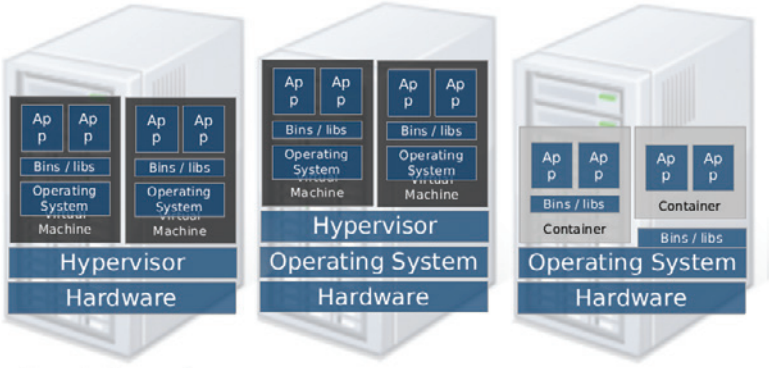

Type 1 Hypervisor | Type 2 Hypervisor | Linux Containers

IBM Global Technology Services

OLSON®

POWER DISTRIBUTION UNITS

- Vertical/Horizontal Mounting
- Combination Units
- Power Monitoring
- Remote Monitoring
- Rated at 13A, 16A and 32A
- Bespoke Units
- Robust Metal Construction
- Availability From Stock
- Next Day Delivery

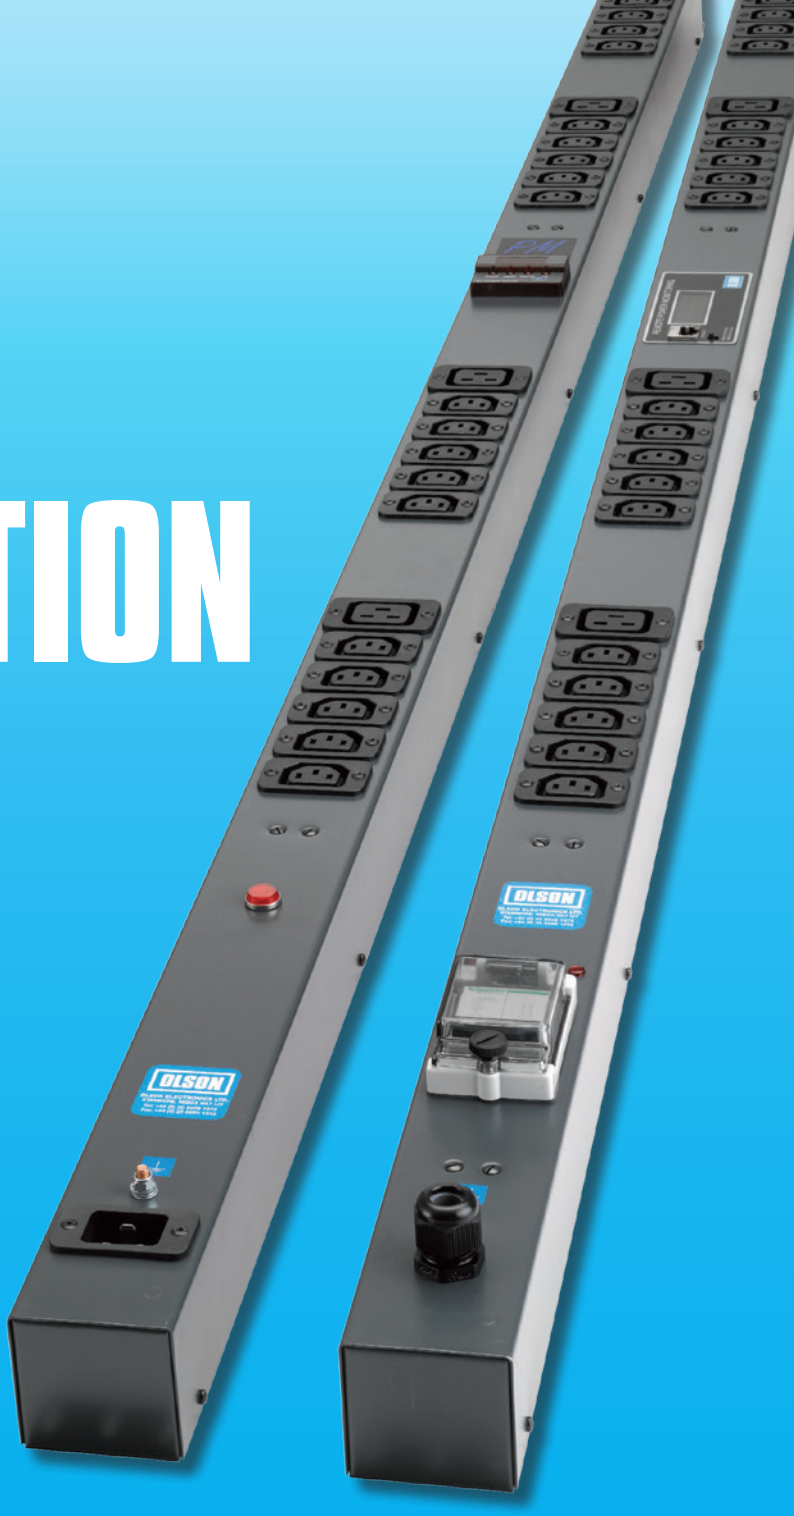

WE'LL BE THERE!

SEE US AT STAND
D342
REGISTER FOR YOUR FREE TICKETS!
21-22 March 2018 ExCeL London

DATA CENTRE WORLD

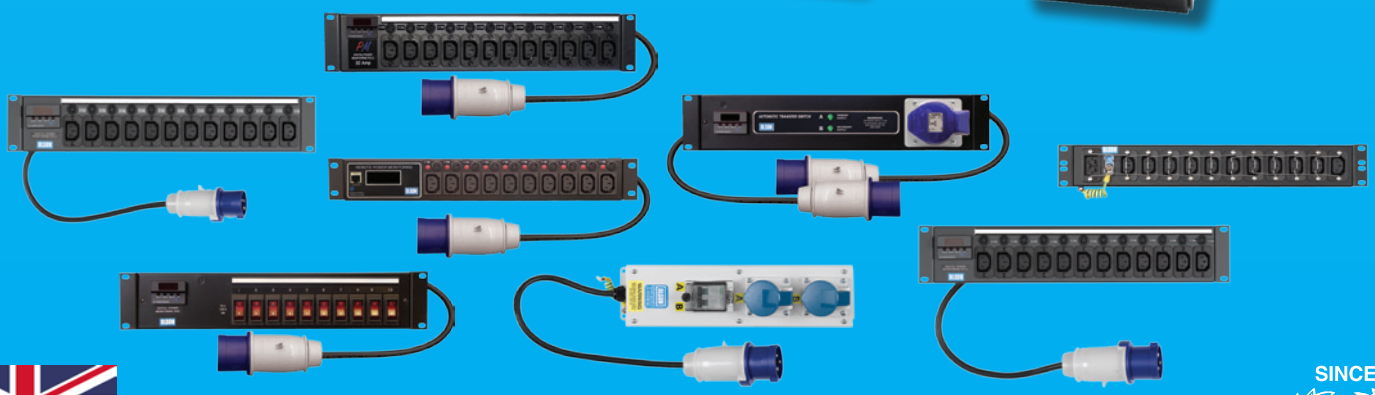

Designed and Manufactured in the UK

+44 (0)20 8905 7273 | sales@olson.co.uk | www.olson.co.uk

SINCE 1961

Whose DCIM data is it anyway?

Cloud-based data center management tools could enable ever more efficient and resilient facilities but the technology also raises some issues around data privacy and security, says *Andrew Donoghue*

Andrew Donoghue
Technology writer and analyst

It is fair to say that data center infrastructure management (DCIM) software hasn't quite lived up to its initial hype.

True, adoption is increasing and operators are less skeptical of the potential of DCIM than in the past. However deployments haven't hit the kind of levels that many suppliers had hoped for. According to a 2016 survey by Intel and Dell, just over half of operators used some form of DCIM at the time.

Long sales cycles, uncertain return on investment (ROI) and complexity have stymied adoption, as has the lack of clear definition of what DCIM software is exactly. The whole notion of DCIM still attracts criticism in some circles.

But not all data center equipment suppliers chose to launch a specific DCIM platform. Power equipment specialist Eaton has on-premises monitoring tools, and partnered with one-time DCIM supplier CA Technologies, but never launched a full DCIM platform.

"We took our time analysing the situation and in hindsight not jumping on the DCIM train hasn't been a bad outcome considering how much investment some of the others have done," said Tatu Valjakka, product owner for connectivity and interface products at Eaton. "There are quite a lot of DCIM companies that have already perished in the market."

While Eaton resisted entering the DCIM space fully, it is one of a small number of suppliers offering a new class of data center management tools that could overcome some of the issues that have held back DCIM adoption to date.

Data center management as a service (DMaaS) – a term coined by 451 Research - essentially moves some of the functions of on-premises DCIM to the cloud. ▶

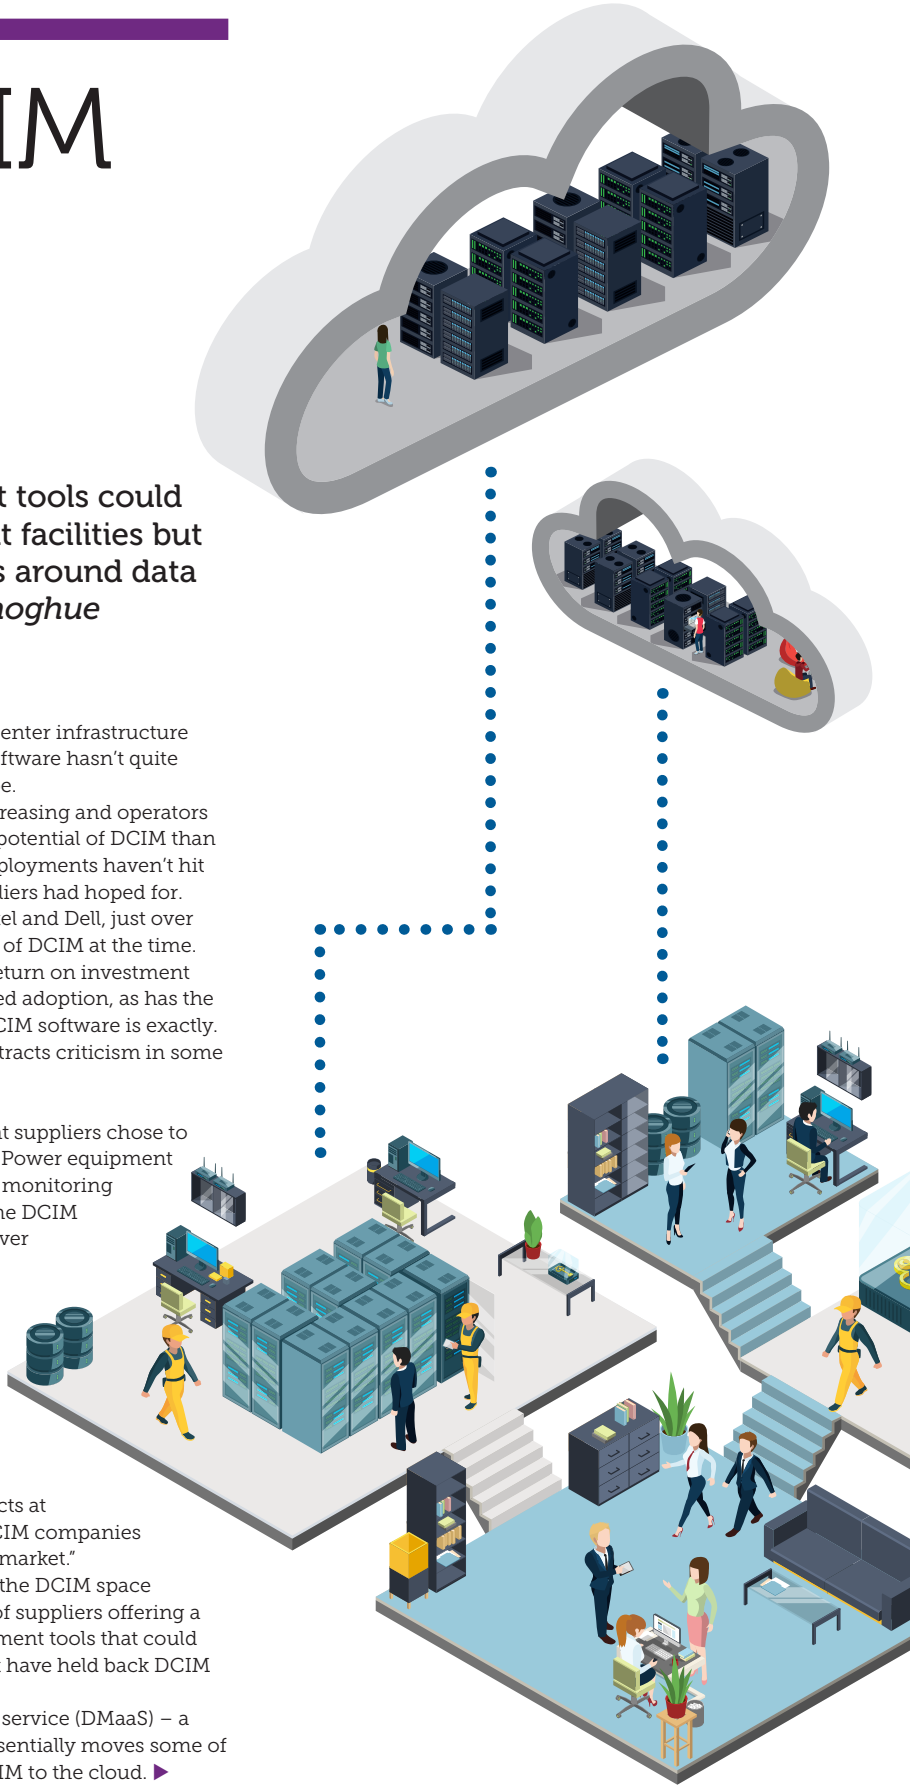

This includes some of the remote equipment monitoring, power and environmental monitoring as well as data management and reporting functions.

The DMaaS tools that have emerged to date will probably be used alongside existing DCIM or building management system (BMS) tools which don't as yet provide capabilities such as IT asset management, an important feature of DCIM, or the kind of control capabilities provided by a BMS.

> *If the customer wants to leave the platform and have their data removed then we will make sure that it is*

However, they do enable a number of additional capabilities made possible by the scale that cloud and big data analytics brings. To date, only a few of the more than 60 suppliers of DCIM have launched DMaaS services but others, such as Vertiv (previously Emerson Network Power), are likely to follow if the initial platforms prove popular.

Eaton launched its DMaaS platform in early 2016. PredictPulse is essentially a cloud-based remote monitoring tool. It is only available in the US to date but enables operators to view real-time power diagnostics, using online dashboards and mobile applications.

Schneider Electric is also one of the early pioneers of DMaaS. It launched its StruxureOn service in 2016 but subsequently rebranded it to EcoStruxure IT in mid-2017.

Essentially, EcoStruxure IT collects data from multi-vendor (not just Schneider) infrastructure equipment and stores it in the cloud. The data is then aggregated in what is described as the 'EcoStruxure data lake' where it is processed and analysed.

As with some on-premises DCIM tools, users are able to receive alerts and other monitoring information – using a smartphone for example – on how equipment in a specific data center is performing.

But the real value-add with DMaaS is that data harvested from all the users of the service – in Schneider's case 1,000 customers with more than 100,000 cross-platform connected devices to date - is aggregated and analysed using big data analytics and machine learning. This enables new capabilities such as a user being able to compare the performance of a facility against global benchmarks.

Another benefit is the ability to use historical trending data, or the data on the performance of multiple pieces of equipment, to more accurately predict when a device may need to be serviced: so-called predictive maintenance. Eaton and Schneider's platforms both provide this capability.

"We are able to predict incidents before they happen," says Schneider's VP and GM of data center software, Kim Povlsen. "We can call the customer up and say 'Hey, you might have a problem tomorrow based on what we have seen week over week at your data center'."

But despite the additional capabilities that DMaaS can bring compared to on-premises DCIM, some operators still want to have dashboards and access to monitoring data locally, says Povlsen. For those customers who want to fully retain control over their infrastructure and not send any data into the cloud, Schneider is continuing to support its StruxureWare for Data Centers platform. "We know that there is a large customer base out there who still prefer to stay on premise," says Povlsen.

Another reason that customers may prefer purely on-premise DCIM to DMaaS is concerns about data privacy and security. However Povlsen argues that the data that Schneider is collecting has limited value to a potential criminal or hacker so customers shouldn't be overly concerned. "We are only monitoring the machine data – UPS systems, PDUs, cooling units - from the critical infrastructure," he says. "So things like temperature and power, that is all we do. So you could say that it is relatively useless data if you don't know what to do with it."

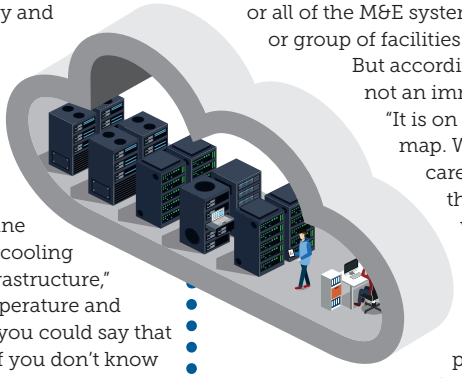

Even so, Povlsen is adamant that the customer retains ownership of the data collected by EcoStruxure IT no matter where it is located. "The data is purely owned by the customer," he says. "If the customer wants to leave the platform and have their data removed then we will make sure that it is."

Instances of customers requesting to have data removed are rare however. "It is something we offer but it is somewhat a rare occurrence but for us it is important to let the customer know that they are in charge," says Povlsen.

Eaton's Valjakka agrees that data security and privacy is not such an issue because of the kind of data being collected. "I would say the data ownership issue is more pressing when you have actual workflow data, people data, operational data. We just collect the internal workings of UPS systems," he says. But, that said, it is customer owned and we will use it for prediction, for developing the models, feeding some of the insight into product development but we are not selling the data to anyone or making money out of the data directly in any way."

But while DMaaS suppliers appear to have a tight rein on issues around customer data, the technology also raises some other data ownership questions. For example, as well as being able to glean useful data on how its own equipment is performing in a customer's site, Schneider is also able to gain similar insights into equipment from rival suppliers.

Since Schneider is able to aggregate data from potentially thousands of pieces of equipment, this could provide insights that even the rival supplier doesn't have access to. "We are thinking about what we can do," says Povlsen. "As of now we don't see any explicit value in sharing this data with our competitors."

The ultimate end game for DMaaS could be for Schneider, Eaton and others to extend cloud-based monitoring capabilities to include remote control of specific equipment or all of the M&E systems in an entire site or group of facilities.

But according to Povlsen it is not an immediate priority. "It is on the far end road map. We want to be very careful about these things. Right now we have a strict policy around one-way data direction in order to provide the proper level of cybersecurity."

Ultimately, DMaaS has a number of advantages over traditional DCIM tools but despite the assurances from suppliers there will be some operators who are just not ready to share data with the cloud. The reality is that DMaaS will make sense for a certain group of users and it will most likely be used alongside traditional DCIM for now.

Over a longer term, however, the idea of a supplier providing full remote monitoring and control of specific equipment, or an entire site, could prove very persuasive for progressive operators - as long as ownership and control over data continues to remain clearly defined. ●

Feeling the heat

Upgrading a data center to get a lower PUE is a challenge, but a PUE of 1.1 in the sub-tropical South African climate is a serious achievement. *Peter Judge* reports on *DCD* Award winner Teraco

South Africa's climate is hostile to running efficient data centers: the high average temperatures limit the amount of cooling that can be provided by outside air, so facilities usually have to resort to alternative, energy-intensive cooling systems.

Colocation provider Teraco has just opened the largest commercial data center in Africa, near Johannesburg. While it was building the Riverfields facility, the company updated its Isando 7 data center, which can now hit a PUE of 1.1 that would be the envy of providers in many a colder climate.

Situated on the highveld plateau, Johannesburg has a cooler climate than much of South Africa, but it still reaches an average daily maximum of 25.6°C (78.1°F) in the summer month of January, and temperatures can go as high as 40°C (104°F).

Founded in 2008, Teraco offers colocation services to the whole of Africa, based on South Africa's superior electrical power and communications infrastructure. Its data center in Isando, Johannesburg, opened in 2010, and has been expanded since.

The site was built with traditional CRAC units for indirect free cooling, and Teraco's upgrade doubled the cooling capacity to 2MW, increasing efficiency without sacrificing reliability. This feat won it the DCD Energy Efficiency Improver's Award, presented at the end of 2017.

To meet its target, Teraco told *DCD,* the only feasible solution was a chilled water system, which delivers free cooling for approximately 65 percent of the year, and provides a steady supply of water at 14°C (57.2°F). The IT space uses cold aisle containment and the waste heat takes the circulating water up to 24°C (75.2°F).

Further cooling could be provided using an adiabatic (evaporative) system, but this would increase the site's environmental impact by using large quantities of water. "No additional water is used except for the 'one-off' filling of the system," Brendan Dysel, head of data center systems at Teraco, said in the project description.

The data center uses three CyberCool 2 chillers and 14 CyberAir 3 chilled water units from climate control vendor Stulz, which consulted with Teraco, helping the provider to implement a system with sequencing control and sensors.

As well as efficiency, Teraco's customers want reliability - and any interruption of cooling could quickly result in facility overheating and cause systems to shut down.

The cooling system is backed with an uninterruptible power supply (UPS), while an external UPS also backs the controller and the pumps.

In the event of a power failure, buffer tanks provide a reserve of chilled water that can keep the data center cooled for ten minutes - more than enough time, since the chillers have a "quick start" feature which will bring them up to 100 percent load within two minutes.

Peter Judge
Global Editor

The three chillers provide N+1 redundancy for the cooling needs of the data center, but in mixed mode, or free-cooling mode, they are all run in parallel. This reduces the overall energy consumption because they operate more efficiently at part load. It also means that the system will remain online if any one of the chillers fails. Only two chillers are used when the site is operating in DX mode.

Any chiller can be taken offline for servicing without shutting down the data center, making the system concurrently maintainable.

"[This] is crucial for reliable uptime," Dysel said. "If you have to shut down the total cooling for maintenance reasons, you have to shut down your data center, which is of course undesirable."

This variety of modes sets strong requirements for the air handling at the site. The chillers provide real-time data and control. Energy meters were installed throughout the data center and showed a measured efficiency within two percent of the predicted figure.

The story is not over. Teraco is monitoring the system through its building management system (BMS) and data center infrastructure management (DCIM) systems, and is working with Stulz to optimize the mixed and free cooling modes further, for example by using sensors installed throughout the server room and chiller plant, and perhaps increasing the water temperature to extend the use of free cooling regime.

"The ultimate goal is to deliver even greater energy efficiency on an ongoing basis," Dysel added. "The DCIM can create a heat map of the server room to identify any hot spots that need attention."

Innovation on ICE

Being a major data center destination, Sweden has a unique research establishment pushing energy efficiency forward, reports *Tanwen Dawn-Hiscox*

Tanwen Dawn-Hiscox
Reporter

Two years ago, Sweden was placed top in the European Innovation Scoreboard by the EU Commission. This vindicated years of educational policy, and reflected a national willingness, both public and private, to invest in research and innovation, and create partnerships joining academia and enterprise to drive the country's economy forward.

Sweden also harbors the perfect conditions for building the energy-greedy, monolithic structures that process the reams of data that keep the world's digital wheels turning.

Thus, it is unsurprising that a publicly funded Swedish research institute in Luleå – both an industrial and an academic hub, the location of two of Facebook's hyperscale facilities and the second biggest contributor to the country's GDP – is combining these two fields of expertise, having launched a fully functioning data center devoted entirely to research and development of new cloud and infrastructure technologies.

RISE SICS North, a subsidiary of applied information and communication technology research institute, was launched specifically to allow Europe's academic institutions and businesses to take on data center related projects in a simulated environment.

The institute's data center, the SICS Infrastructure and Cloud Research & Test environment (ICE), was built as a sort of innovation incubator, to boost European capability within the data center industry.

The initiative is based on the premise that the world is set to enter a new era, with compute power in huge demand. To enable this, the data center industry must build on existing knowledge to develop elaborate, software-controlled IoT systems, capable of operating autonomously while leveraging machine learning to its full potential.

The head of research for Ericsson, Tor Björn Minde, is the project's CEO. He told us that ICE was launched to attract companies and projects to the Node Pole region around Luleå. "This will enable a sustainable competence development," he said. "New innovations and skilled people are needed to develop the Swedish economy."

Up from only three staff members in January last year, the facility now operates with twelve people, he said, "and we have a €2.3 million ($2.84m) annual turnaround."

Funding comes from all directions: "From companies - Vattenfall, ABB, Ericsson, Facebook. We get regional and national subsidies, and we have one EU project at the moment as well."

When Minde gave us a virtual tour of his facility, he proudly showed off the DCD Award for Best Data Center Initiative of the Year that the institute received at the ceremony held in December last year. "I bought the cabinet in September and had only one award to put in it. Everybody laughed. Nobody is laughing anymore."

> The facility combines Sweden's penchant for research with its role as a data center location

The 600 square meter facility is a gift that just keeps giving; all manner of experiments and novel equipment fill the room, such as a liquid cooling tank designed and built in-house, a miniature network operations center to test out DCIM systems, and an ABB microgrid controller.

The research data center's first room-in-room module came online in February 2016, based on Dell SmartEdge Rx730xd servers with GPU acceleration and both OpenStack and Hadoop storage.

Then in May 2017, a second, more flexible "lab-like" module was launched, with more power and faster networking connections than its predecessor, and this time equipped with a combination of Dell and OCP v2 servers, standing on raised floors and partitioned using industrial PVC curtains. The latter, Minde tells us, like everything else in there, is something they are testing out. Another setup, more experimental still, will be an OCP lab, fitted with edge PDUs; though for now, only a switchboard has been installed.

Another object of study was brought to the lab by Jon Summers, a senior lecturer at the UK's University of Leeds, also RISE SICS North's scientific leader and an expert in CFD and energy efficient cooling systems: a glass box the size of a room, designed by DigiPlex, which is in fact a small, contained wind tunnel for airflow experiments.

Researchers are also exploring the possibility of using chimneys to create drafts powerful enough to action fans in HVAC systems, in the form of four retrofitted 19-inch racks containing OCP servers, placed in a cross to allow for a chimney in the center.

Soon, Minde tells us, the site will have its own solar panels, a thermal storage bank and a heat pump for waste heat reuse.

Throughout the facility, networking operations are monitored by Zabbix 3.0 open source software, and additional sensors have been placed on PDUs, power switches, servers, network switches, UPS systems, coolers, power meters, thermometers and corrosion sensors, collecting between nine and ten gigabytes a day from 30,000 measurement points per module. The data is then fed to a Hadoop cluster using Apache Kafka processing software.

Projects currently underway seek to improve knowledge in the fields of data center automation, air-flow technologies, "industrial symbiosis," as Minde calls it, or micro-grid integration, data collection and monitoring, and machine learning for the data center.

The study of "Artificial Intelligence for Datacenter Microgrid Interaction," funded by the Swedish Energy Research Centre, looks at the possible applications, as well as the pros and cons, of integrating AI into a single controllable entity that would link a microgrid with a data center, in view of making it possible to take temperature, solar radiation, electricity cost and workload forecasts into account when operating an edge module.

"DC2BD" is a heat reuse project backed by Forest Developing Technologies and Future Eco, which is exploring the possibilities of using heat from data centers to dry biomass.

More widely, it seeks to demonstrate that creating a circular economy, where a data center would help create bio-coal or biofuel, and generate heat as part of a district heating system, for instance, is possible.

The facility is also the basis for a data analytics project called D-ICE, a platform that enables data owners and data scientists to share information - the former gaining from in-depth analysis of their data using new tools, and the latter obtaining access to the reams of information necessary to test out new methods.

The 60 plus companies participating in projects (although, Minde tells us, approximately 40 are actively involved) are continuously invited to see the research institute and improve their own efficiency and sustainability practices, while contributing to the wider European market.

For Minde, innovation goes in cycles, from the bottom, through the hardware stack, up the software layers, in alternating phases.

Explaining the purpose of the institute last year, he said: "We [at RISE SICS North] improve on the foundation stuff; maybe now we're going into a new hardware cycle; we have a lot of new innovations in the hardware side, there's a lot of interest in the new types of hardware, and then the software will innovate on top of that. We at SICS are an integral part of that development."

As for the future plans, he tells us, "in the next two to three years, we hope to double the [institute's] activities and build a new facility that is two to three times larger," making it between 2,500 and 3,000 square meters - a much bigger space in which to test out the technologies of tomorrow.

Battle for the South

Virginia Toledo
Editor LATAM

Argentina wants to win business from Chile, say *Virginia Toledo, Sebastian Moss and Peter Judge*

Argentina's pitch

- Cool climate
- Crackdown on corruption
- Argentina's Investment and International Trade Agency (AAICI)
- New data protection legislation
- Lowered import taxes on servers
- Dark fiber available near Buenos Aires

Growing demand for data processing has come to Latin America just as it has to the rest of the world, and countries are competing to become the local hub for colocation and cloud infrastructure.

South of the Tropic of Capricorn, in the region known as the Southern Cone, Chile has been far ahead of its rivals, Argentina and Uruguay, attracting numerous data centers, including cloud facilities for Google and IBM. But two years ago, President Mauricio Macri took charge in Argentina, with the intention of attracting data center projects along with other sources of foreign investment - and his first targets include Amazon Web Services.

After AWS expressed an interest in launching a data center in Chile, Macri met with Elaine Feeney - Amazon's VP for global infrastructure expansion - in New York in November last year. Amazon also visited Argentina to assess the market for cloud services and the local tax regime.

But Chile's President Michelle Bachelet traveled to Amazon's headquarters in June. AWS then signed a memorandum of understanding with the government of Chile "to modernize government services within the country," and a government spokesperson went on record to say AWS is interested in placing a facility in Patagonia - on the southern tip of the continent.

As well as getting close to customers in this part of the continent, AWS is interested in the cool climate which makes data centers more efficient, and also in Chile's economic and political stability. The country hopes to have a new 20,000 kilometer submarine cable linking Patagonia and China, a $550-$650 million project that would be the first direct cable from Asia to Latin America.

As for Argentina, Macri can point to the country's efforts to end years of financial and political unrest, with a crackdown on corruption that has won business in the financial and petrochemical sectors.

One year ago, Argentina's new Investment and International Trade Agency (AAICI) kicked off a campaign to promote the southern region of the province of Buenos Aires as a hub for data centers. The project is coordinated by Andrés Tahta, EVP for renewable energy, telecommunications and tourism at the AAICI.

In 2016, an AAICI investment forum gathered 1,000 CEOs, 400 government representatives and national and international leaders to show the world what Argentina has to offer. "At this forum we met many investors," Tahta told *DCD* in June 2017. The department has since been meeting multiple data center companies at global, regional and local levels, he said.

Points in favor of Argentina include new data protection legislation and a new regime of import taxes, both of which are in progress, along with the cool climate and the fact that there is little risk of floods or earthquakes. While Chile awaits its direct link to Asia, Argentina can point to the availability of existing fiber connectivity across the Atlantic and Pacific.

In southern Buenos Aires, there is plenty of dark fiber available, Tahta said: "We are talking not only of the fiber of the leading telcos but also 33,000km belonging to state-owned telco Arsat." So far, 22,000km of Arsat's fiber has been lit and the rest is being illuminated bit by bit. Meanwhile, Seaborn plans to open a submarine cable linking Brazil, Argentina and the US in 2018.

Tahta told *DCD* his work is being supported by other ministries including communications, justice, and the Undersecretary for Renewable Energy.

Argentina will need to modify its data protection law, which dates from 2000. The new legislation will "strike a balance between regulations from the US and Europe," he promised.

The new import scheme will be very friendly to data centers, he said, with tariffs on imported servers and other data center components, "at or near zero, for as long as possible." He hopes the idea would be applied to subsequent upgrades as well as the initial investment: "We are working with the Ministry of Commerce to eliminate selected import duties. We are also exploring the possibility of the creation of tariff benefits for data centers."

> *Many of the companies we are talking to are in the southern region of Buenos Aires*

All data centers need reliable power, and increasingly prefer that power to be from renewable sources. Argentina has made great efforts over the past year and a half to increase the availability of alternative sources of energy, Tahta said: "A draft invitation to tender sought 1GW power within an overall plan to reach 10GW in the next ten years."

In the end, the the initial number was raised to 2.4GW, he said. From having just 1.9 percent of renewable power in 2016, the country has reached nine percent, and hopes to have 20 percent by 2025.

Argentina is developing solar energy and harnessing wind power, which is plentiful in the southern region of Buenos Aires. The government plans to approve new rules that will enable renewable energy producers to feed it into the grid and make it available to data centers.

"Many of the companies we are talking to are in the southern region of the province of Buenos Aires," he said. "When we told them that each data center requires approximately 100- 120MW, several told us that they were considering putting windmills higher than those built under government contract.

Although this is a national project supporting data centers installed anywhere in Argentina, Tahta insisted that the ideal location is the southern province of Buenos Aires, which is much cooler than the capital itself, is flat and not subject to flooding, and has wind farms less than 100km away.

Also, the Las Toninas cable landing station, less than 300km away, provides low latency links to important cities including São Paulo, Rio de Janeiro, Lima and Santiago de Chile: "This is critical because the project aims to reach all Latin America, not just Argentina," Tahta said.

Compared to Argentina, Chile has higher seismic risk, and is "isolated and far from the main centers of data consumption," he said. Tahta also compared Argentina to Brazil, which has plenty of renewable energy, but also has a hotter climate. He implied that Brazil is lagging behind on data center-friendly legislation, warning that investments there could get caught in a "tangle of bureaucracy."

Since *DCD* spoke to Tahta, the initiative has shown results: US edge specialist EdgeConnex has plans for a 10MW facility in Buenos Aires. Francisco Cabrera, Minister of Production, said: "The new Buenos Aires EDC will provide the region economic growth... making Argentina a more attractive location for data investment by enterprises across a wide range of industries."

While a feather in Argentina's cap, the project isn't necessarily a win over Chile: EdgeConneX also has plans for a facility in Chile's capital, Santiago.

There are suggestions that AWS may do the same: the cloud giant is expanding so fast it too may need capacity in both countries. ●

Additional reporting by Sebastian Moss and Peter Judge

Edge takes a central role in APAC

As billions of people come online, the region needs to build more data centers, move to the cloud, and focus on edge services, says *Paul Mah*

Paul Mah
SEA Correspondent

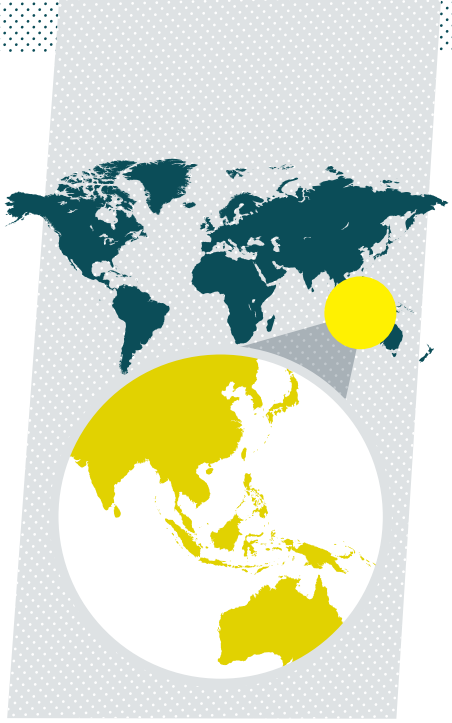

There is no doubt that there is huge potential in the Asia Pacific (APAC) region. It has an enormous population base of 4.5 billion people, and yet slightly less than half are connected to the Internet.

It should therefore come as no surprise that it is the fastest growing region for data centers in the world, according to Structure Research, which also projects that APAC will become the biggest data center market in the world by 2020.

As for 2018, it will likely see widespread adoption of hybrid cloud in the region, Heng Wai Mun, the executive director of OneAsia Network, told *DCD*. He sees the cloud as an enabler for innovation, and thinks that late-starters making the move towards digital transformation will embrace it in some form or other.

"The conversation over the last few years has always been on which option [public cloud or on-premise] is better, but it will switch this year to how to get the best of both worlds. The hybrid approach will provide organizations with an agile, scalable platform that can support the business even during times of rapid change and growth," noted Omer Wilson, the APAC senior director of marketing at Digital Realty.

"We certainly see more businesses adopting the hybrid cloud infrastructure moving forward. As businesses mature in their use of new technology such as artificial intelligence, predictive analytics and IoT, they are beginning to realize the benefits of on-premise solutions," Wilson said, though he concedes that the cloud remains the most viable option when it comes to scaling.

Glen Duncan, a senior research manager for infrastructure sector at IDC, agrees that organizations are leaning towards hybrid cloud deployments in Southeast Asia: "It's generally agreed that the data center environment for the foreseeable future will be a multi-cloud one. This means that enterprises will utilize a mix of on-premise, third-party, co-location, private cloud, hosted cloud and public cloud."

Duncan was quick to note that adoption decisions will inevitably be impacted by a variety of factors. These can include the nature of workloads, legacy decisions made by the team, budget sizes, technology maturity within the organization and cultural factors.

The cloud is not a one-size-fits-all platform, explained Philbert Shih of Structure Research, and organizations will make an independent evaluation to identify the most optimal mix of cloud services and outsourced infrastructure for themselves.

Shih told *DCD*: "What combination of cloud services minimizes cost and maximizes my performance and efficiency? These are the questions end-users are asking and it means that hybrid deployment models are going to be the norm as the world increasingly outsources."

Deployment choice ultimately boils down to a matter of wearing the shoe that fits. Echoing the conclusions of a recent study by 451 Research, Duncan highlighted a growing trend where organizations are moving some types of workloads away from public cloud environments, typically those that are no longer cost-effective due to their scale.

This meant that on-premises systems such as hyper-converged systems should not be discounted. He said: "On the IT infrastructure front, appetite for hyper-converged appliances will grow in [the region], although current adoption is still low. Drivers for growth of these appliances include ease of deployment [or] management, cloud-readiness and suitability for edge environments."

Indeed, Wilson thinks that colocation will continue to maintain a crucial role even as IoT and edge data center deployments become increasingly prevalent. Citing a Gartner report from last year that spotted a massive shift to hybrid infrastructure services, he suggested that 2018 will see businesses increase their reliance on colocation services to optimize cost and boost efficiency.

The result? Forward-thinking data center providers will probably blend cutting-edge technologies such as artificial intelligence with data center infrastructure management (DCIM) tools to increase energy efficiency, improve uptime and reduce operation costs, he said.

If there is one common topic mentioned by every analyst, it is the growth of edge computing. Heng from OneAsia thinks that edge computing will move from prediction to reality this year, while Duncan described it as a "big shift" from a centralized data center architecture towards a decentralized one.

"[The] multi-cloud environment is not just on-premise, colocation and cloud in all its flavors. It pushes out beyond this to include a plethora of devices and sensors at the edge to be served by these edge data centers," he said.

Duncan noted that such facilities will also adopt smarter automation and management technologies, primarily due to the lack of human intervention available at edge sites.

What could be the impact of edge computing, as it takes a more central role in the region? Jabez Tan, research director at Structure Research, thinks the gradual shift to the edge, coupled with considerations such as data location requirements, interconnection and ecosystem density, will move the needle towards markets in places such as Jakarta, Seoul, Taipei and Cyberjaya in Malaysia.

"We believe there will be a gradual and incremental deployment strategy in the tier 2 markets with the majority of new data center deployments still housed in the top Asia Pacific markets, driven by interconnection and ecosystem density," Tan said.

And edge computing or not, there is no question that the primacy of data is here to stay. Considerations such as fast connectivity, accessible solutions marketplaces and access to real-time data are as important as a robust, energy-efficient hosting facility.

"It is not just about providing the best connectivity or uptime; it will need to offer a complete ecosystem that businesses can rely on for future growth," he summed up. ●

Bison dollars

> *"Every Bison dollar will be worth five British pounds... once I've kidnapped their Queen."*
> General M. Bison in *Street Fighter*

The world has gone mad for cryptocurrencies. As I'm writing this, the price of bitcoin has started to recover, following a crash from more than $19,000 in mid-December to below $8,000 in the beginning of February.

The reasons given for this are many and varied, but the truth of the matter is nobody knows for certain. Now, if a traditional asset would show this level of volatility, no serious investor would touch the stuff. But bitcoin is a technology, and in the famous words of Arthur C. Clarke: "Any sufficiently advanced technology is indistinguishable from magic."

Computers have transformed the world; the pace of innovation is accelerating, and people are more likely to believe in something that would have sounded like a fever dream just a few years ago. If we have quantum computing and rockets that send cars to space, why not magical Internet tokens that can make anyone rich?

Full disclosure: after first hearing about bitcoin in 2012, I became a passionate advocate for cryptocurrencies. Back then, it was about clever mathematics, open source and digital anarchism. Bitcoin was going to take down the banks, liberate the people from the yoke of the government - and it had a cool origin story.

Instead, it became a tool for crime and deception. It's not just about buying drugs online anymore: the recent ransomware epidemic was enabled by wide availability of cryptocurrencies, transactions which supposedly cannot be tracked (they can, but the process is complicated and expensive).

Recent strains of malware no longer seek to harvest personal data – instead, they repurpose unsecured devices for crypto-mining, once again making it easier for cyber criminals to be paid. Cryptocurrencies have enabled countries like North Korea and organizations like ISIS to circumvent regulated financial systems. And do I even need to mention the environmental impact of industrial-scale mining?

And yet, clueless bedroom investors continue to think it's a great idea to sink their savings into that thing they heard about on the Internet. One of the reasons cryptocurrencies have spread like wildfire is the power of conversion: after buying their first bitcoin, owners turn into evangelists, since every bitcoin purchased by their friends, colleagues and social media contacts increases demand, supposedly driving the price upwards and increasing the value of bitcoins already owned. That's the basic economic model, right?

Cryptocurrencies are far from harmless. Friends don't let friends invest in bitcoin.

Max 'Coinless' Smolaks
News Editor

You can tell Max he's wrong about bitcoin by email (max.smolaks@datacenterdynamics.com) or on Twitter (@maxsmolax)

Thank you!

Chain of Hope would like to thank guests at **the DCD Global Awards 2017** for mending little hearts. An amazing £53,395 was raised on the night for children with heart disease!

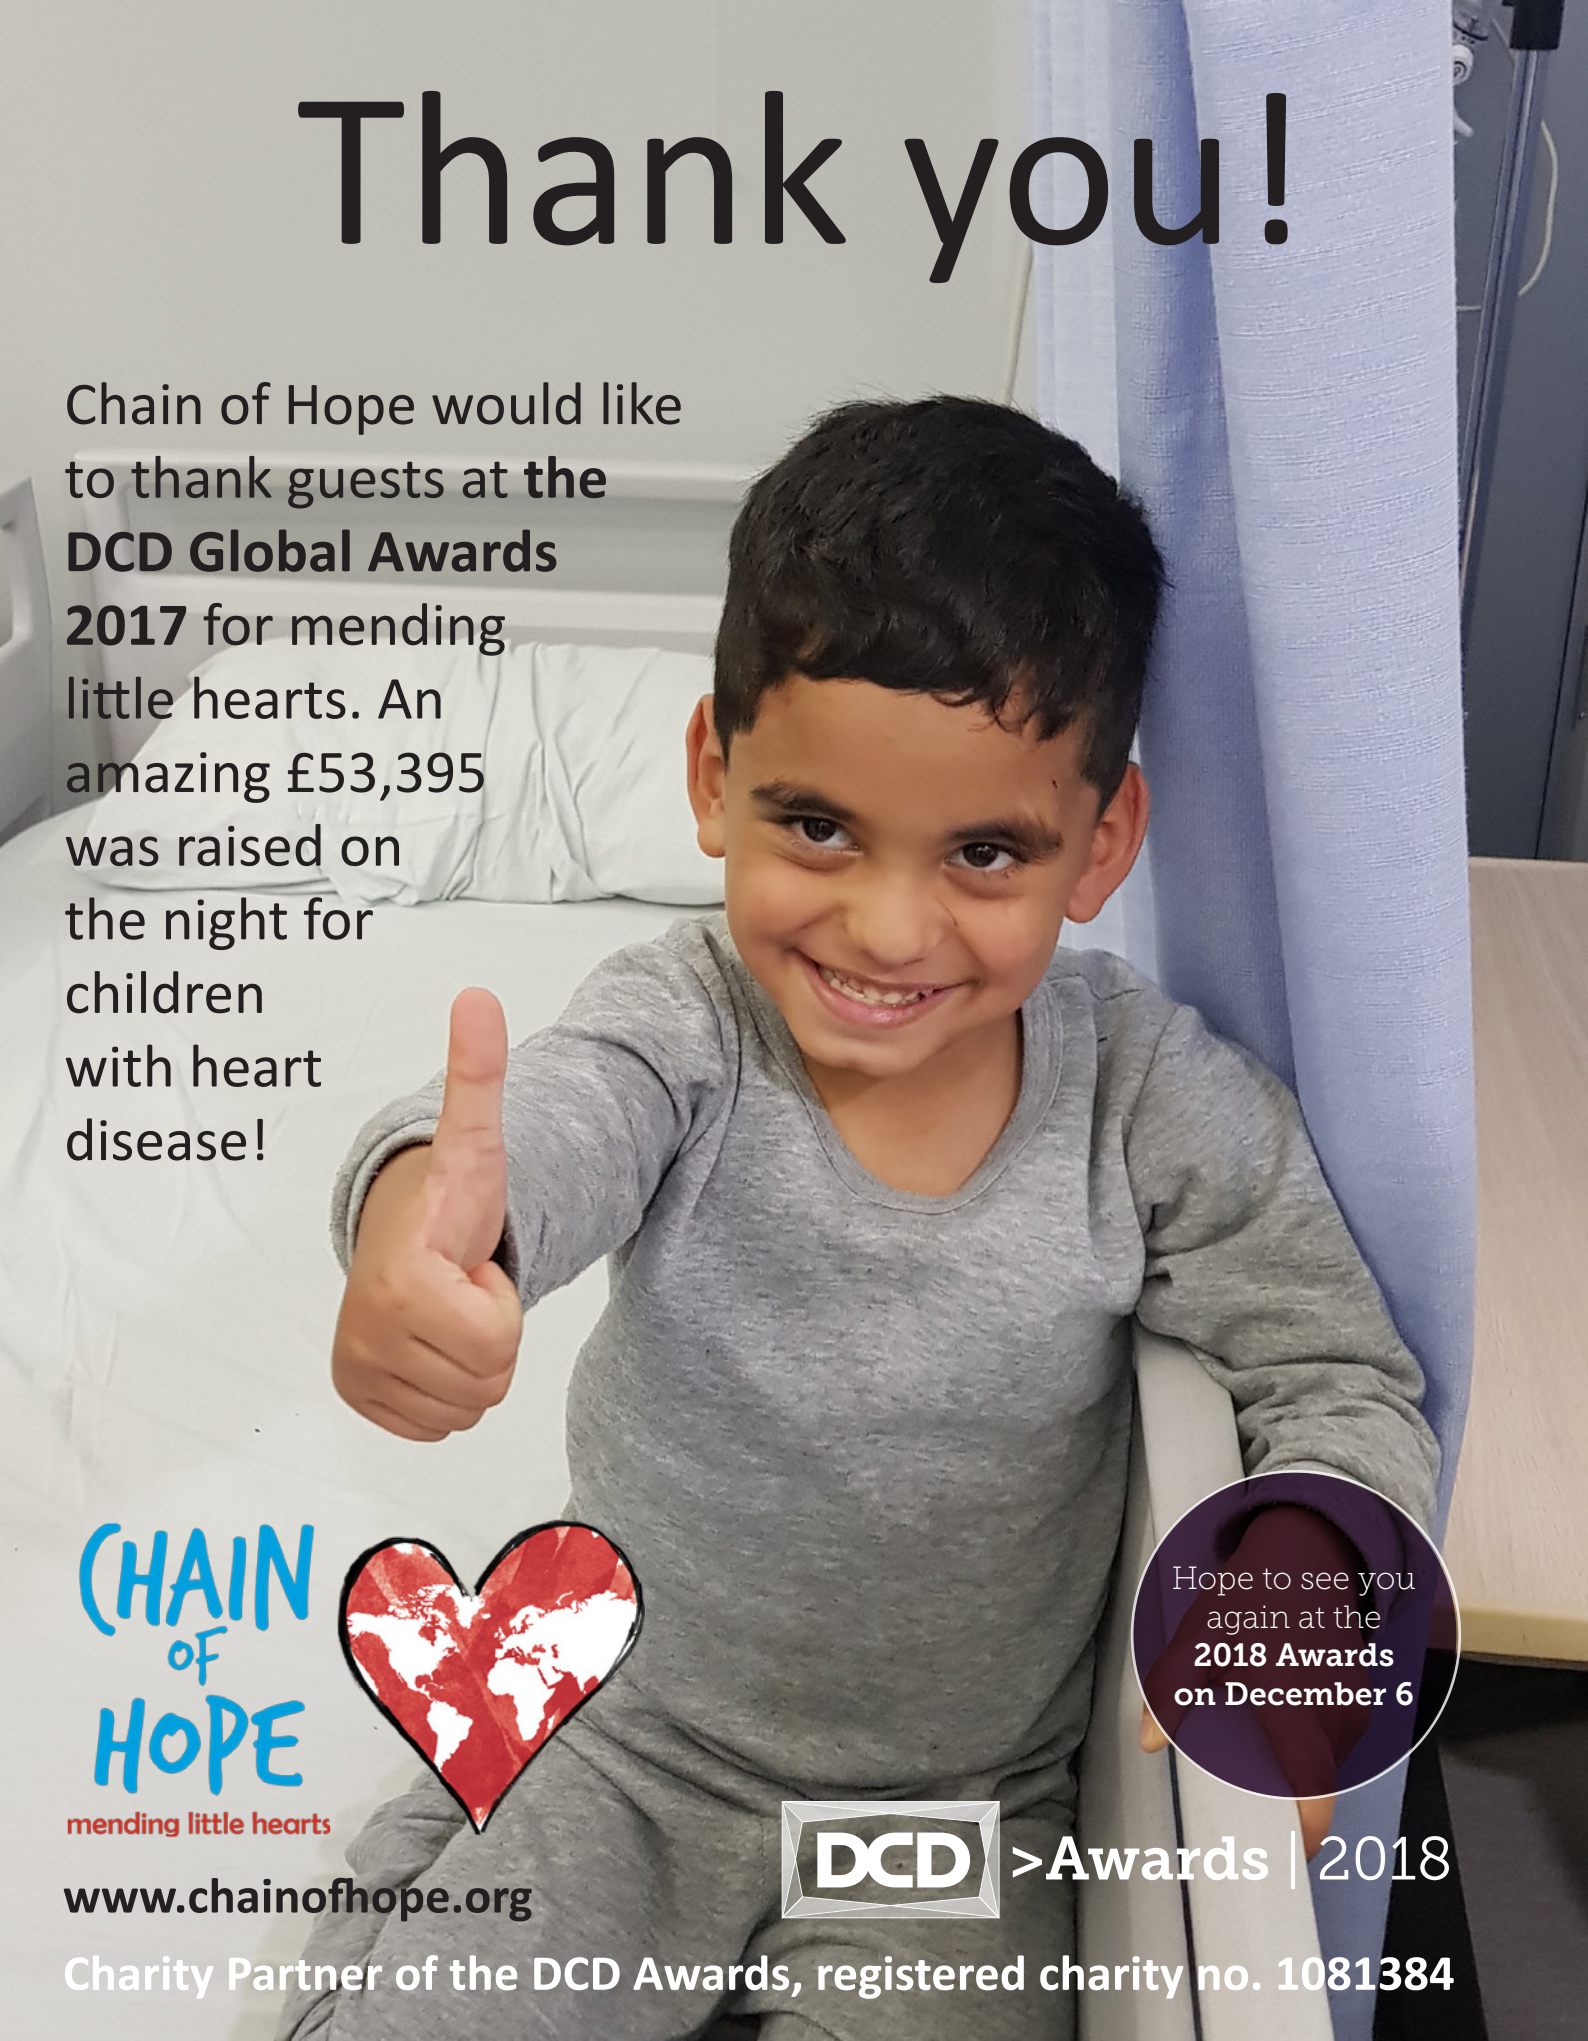

CHAIN OF HOPE
mending little hearts

Hope to see you again at the **2018 Awards on December 6**

DCD >Awards | 2018

www.chainofhope.org

Charity Partner of the DCD Awards, registered charity no. 1081384

CORNING

Simple and new meets tried and true

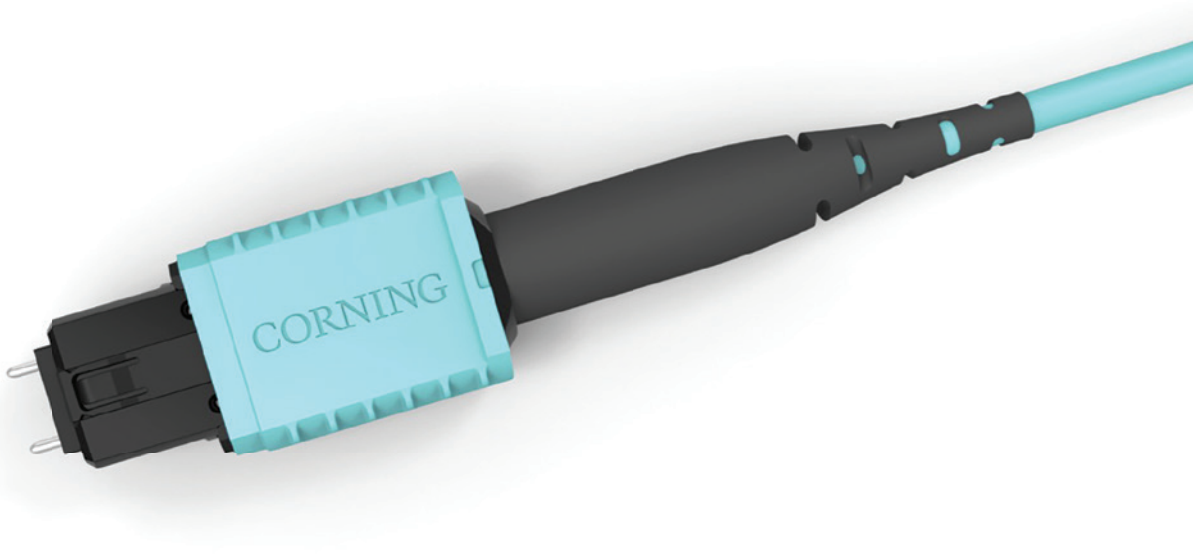

Our industry-leading MTP® connector family has a new addition – the MTP® PRO connector. Now, with no special skills or training required, you can easily perform polarity and pin reconfiguration in the field without removing the connector housing or exposing fibers.

The new MTP PRO connector offers simplified and robust field configurability, ease of use, and enhanced performance combined with the MTP connector you know and trust.

Visit **www.corning.com/mtp-pro/dcd** to learn more about the features of the MTP PRO connector.