April/May 2017
datacenterdynamics.com

Why the human factor is still the central issue

DATA CENTERS OF THE WORLD

POWERED BY PEOPLE

10 beautiful data centers

Here are some facilities that manage to get their work done, and still impress us with their looks

Finding new frontiers

Building a data center is hard enough. Building one in Angola is even tougher

Re-shaping servers

They looked solid and unchanging, but now servers are being remolded by the creative minds of the industry

Contents

April/May 2017

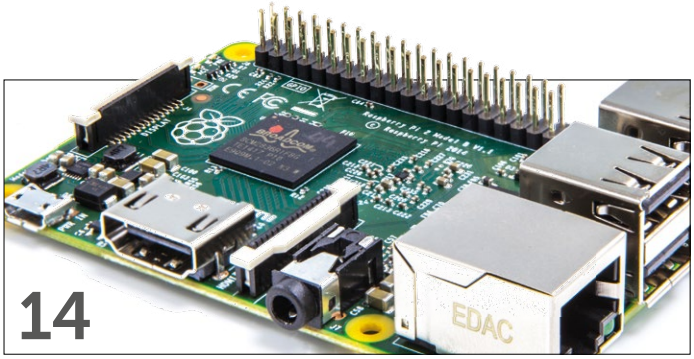

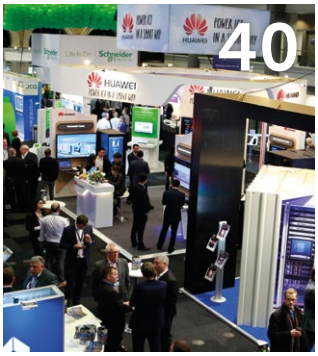

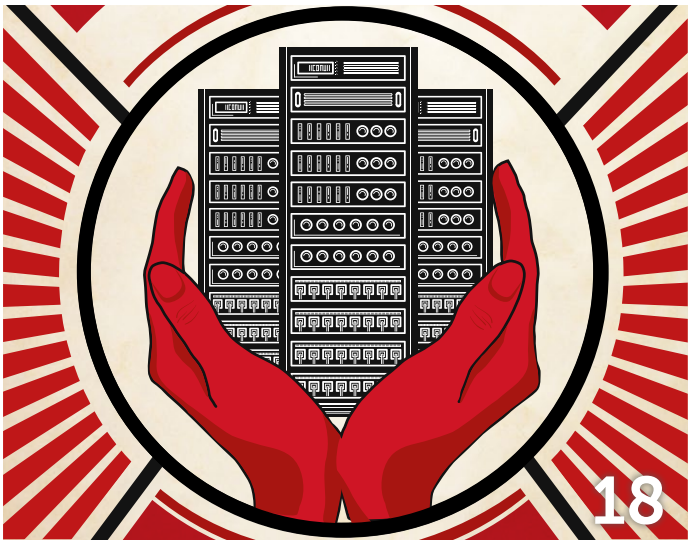

From the Editor

Let's hear it for the humans!

Amazon's S3 storage service went down for several hours early this year, taking large parts of the public-facing Internet with it. The cause was given as "human error."

But what is this "human error" we talk about? An Amazon staffer mistyped a command and deleted virtual servers handling index files. Amazon has now set safe limits to how much capacity a staff member can remove at one go. I'd say the human error was in the original design of the system.

In data centers, we collaborate with ever-more-complex and intelligent hardware. This month (p18) *DCD* puts the spotlight on this symbiotic relationship.

We want data centers to reflect the heart and soul of the people that build and operate them

Could beauty help humans live and work in data centers? A good working environment could help keep staff motivated to handle unprecedented change.

This month we feature ten of the most beautiful data centers we know (p30). This year's DCD Awards will include a special prize for data center design, and this month's roundup is a first step to a short list, so tell us your favorites.

We aren't looking just for architectural excellence. We want data centers whose physical features reflect the heart and soul of the people who build them.

Software is less visible, but it's changing things, as software-defined data centers (SDDC) come closer to reality.

We can build automated services based on pools of compute, storage and networks - see the SDDC supplement that ships with this magazine. The next step is to match these services to business requirements, and deliver what might become known as "business-defined data centers."

Hardware is not standing still meanwhile. After many years of standardization, different shapes and designs of servers are emerging (p37), says Dan Robinson.

And networks have become a crucial tool for data centers, as colocation providers bundle connections to make their own ecosystem stand out amongst rivals. Martin Courtney (p24) finds how these providers are tying down the cloud.

Builders battle spiders, floods, power cuts and armed raiders to deliver data centers to emerging markets (p27).

We believe that these facilities get more users online, and help drive human growth in the process.

22%

of power outages are caused by human error (Ponemon/Vertiv study, 2016)

Peter Judge
DCD Global Editor

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

MEET THE TEAM

Peter Judge
Global Editor
@Judgecorp

Max Smolaks
News Editor
@MaxSmolax

Sebastian Moss
Reporter
@SebMoss

Tanwen Dawn-Hiscox
Reporter
@Tanwendh

David Chernicoff
US Correspondent
@DavidChernicoff

Virginia Toledo
Editor LATAM
@DCDNoticias

Celia Villarrubia
Assistant Editor LATAM
@DCDNoticias

Paul Mah
SEA Correspondent
@PaulMah

Tatiane Aquim
Brazil Correspondent
@DCDFocuspt

DESIGN

Chris Perrins
Head of Design

Fay Marney
Designer

Holly Tillier
Designer

ADVERTISING

Yash Puwar
Head of Sales

Aiden Powell
Global Account Manager

FIND US ONLINE
datacenterdynamics.com datacenterdynamics.es datacenterdynamics.com.br twitter.com/DCDnews | Join DatacenterDynamics Global Discussion group at *linkedin.com*

SUBSCRIPTIONS
datacenterdynamics.com/magazine

TO EMAIL ONE OF OUR TEAM
firstname.surname@datacenterdynamics.com

Packing more power into the same flexible design.

Introducing **1200 Amp Starline Track Busway**, the newest addition to the T5 Series.

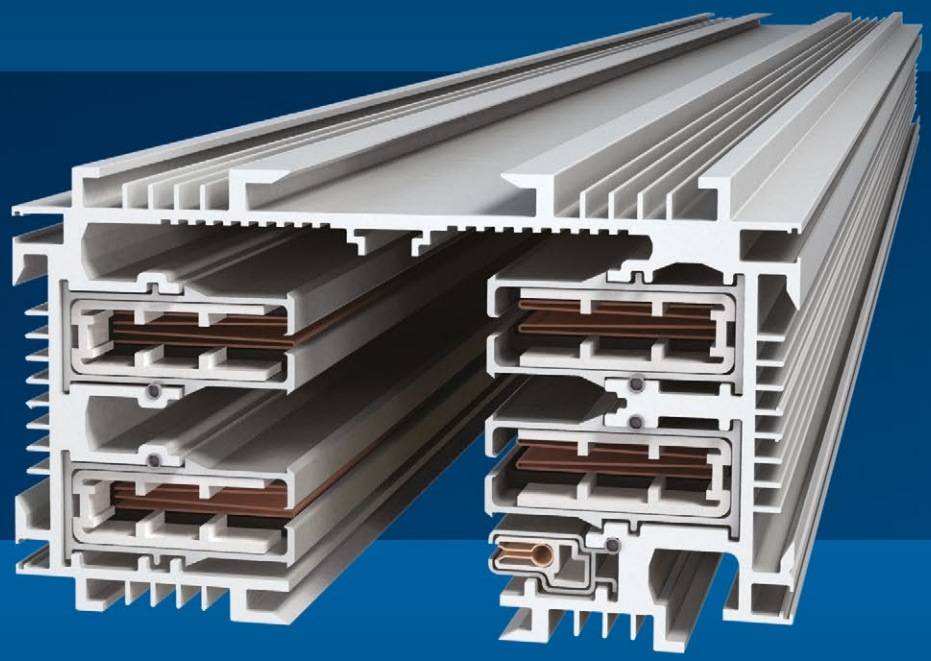

Starline's newest addition to our line of track busway options delivers even more power, without losing any of the flexibility that you've come to rely on from our T5 Series. The 1200 amp busway still has our unique, modular, build-as-you-need design, and is equipped for high-density electrical distribution under the most demanding conditions in data center and mission critical environments. Best of all, the same plug-in units can be used for all systems in the T5 series—from our 250 amp to the new 1200 amp. To learn more about our newest and largest busway choice, visit **StarlinePower.com**.

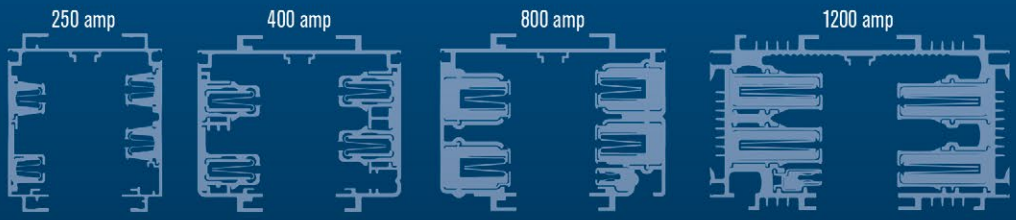

STARLINE®
TRACK BUSWAY

LG U+

Most Eye-catching Datacenter in Asia

Korea, Fastest growing IT hub, becoming a hot place for major global IT businesses

Pack'em up Tight!

Unlimited higher rack power density, for lower Total Cost of Ownership

- The only Carrier Neutral datacenter operated by Tier I Network service provider
- Certified as 'Tier III' by Uptime Institute
- Asia's largest capacity of white space and power (165MW)
- Green datacenter with advanced energy efficiency (air-side economization)

Uptime Institute

Information Security Management

Best IDC Service Provider

For further information ▶ Please refer to our website **globalidc.uplus.co.kr**

Video clip ▶Please search for 'LG U+ Mega Center' at the Youtube YouTube or QR Link

NTT to develop data centers for connected cars

Toyota and NTT have struck a collaboration agreement to develop global technology infrastructure for connected cars.

Schneider launches Micro Data Center (DC) Xpress

At DCD Enterprise New York, Schneider Electric introduced a micro data center range which ships in two to three weeks, with a variety of IT equipment pre-installed.

Microsoft donates $465m worth of free cloud services in a year

Microsoft says it provided $465 million worth of free services to 71,000 non-profit organizations and 350 research universities in the past year via its Microsoft Philanthropies charitable arm.

Human error knocks out Amazon Web Services

In late February, an AWS team member mistyped a command when debugging the S3 billing process, and accidentally removed crucial subsystems, causing internal - and external - chaos. 33 of AWS's own services were impacted, along with countless cloud-based applications and websites.

Vantage bought by Digital Bridge-led investment consortium

West Coast wholesale colo provider Vantage Data Centers has been bought by a consortium of three investors.

Infrastructure company Digital Bridge led the acquisition of Silver Lake's company, with backing from two pension funds: Public Sector Pension Investment Board (PSP) and TIAA Investments.

"Vantage will be the wholesale colocation brand of Digital Bridge," Vantage CEO Sureel Choksi told *DCD*. Digital Bridge already owns a data center company, but Dallas-based DataBank is focused elsewhere.

Vantage is a wholesale provider in Tier 1 markets, whose customers normally use 1MW or more at a time, filling a suite or a whole floor, while DataBank deals with smaller customers, said Choksi.

"DataBank is a retail-focused colocation and interconnection managed services provider, operating largely in Tier 2 markets," he said. "DataBank's customer base typically deploys a small footprint in retail colocation space."

Vantage has four data centers on its flagship campus in Santa Clara, Silicon Valley, with a total of 51MW, and extensive building plans, which will continue unaffected, Choksi told *DCD*.

"There will be no change to our current development plans," he said. "We will continue to build out as fast as we can due to customer demand. This acquisition gives us the opportunity to expand into new markets, given the deep pockets of three investors, who have approximately a trillion dollars under management."

Choksi will remain in his post, while former Digital Realty CEO Mike Foust, an advisor to Digital Bridge, will join the Vantage board of directors, along with Raul Martynek of Digital Bridge.

http://bit.ly/2nhhK9E

Vox Box

Dean Nelson
Head of Infrastructure
Uber Compute

How do we solve the talent gap?

The amount of growth in this industry is huge, and we're looking for people with all kinds of different experience. But the amount of people coming in from universities is shrinking. Students are going into business, they are going into computer science, and they are going into programming. We've been looking for people with experience in different industries, and bringing them in to look at our problems.

http://bit.ly/2oxJmsr

Peter Panfil
Vice President,
Global Power
Vertiv

Are lithium-ion batteries really safe and viable?

A lot of you are thinking "I don't even want to touch Li-ion batteries in my data center! I've heard about Galaxy Notes and hoverboards catching on fire." It's a different tech. It's a chemistry that gives you higher energy density. If you have space or weight distribution issues, Li-ion might be for you. The initial cost is higher than for a VRLA battery, but the CO is compelling.

http://bit.ly/2nLwpOj

Facebook officially announces massive Sarpy, Nebraska data center

In the final weeks of 2016, The Papillion City Council approved a major data center project in Sarpy County, Nebraska on a 146 acre lot.

But the identity of the data center's owner was kept a secret - or at least it was meant to be, with *DCD* finding a paper trail outing Facebook as the company behind the scenes. Now, the social media giant has finally come out of the shadows.

At a press conference at Papillion City Hall attended by Governor Pete Ricketts, State Senator John Murante, Papillion Mayor David Black and Tom Furlong, VP of infrastructure for Facebook, the company officially revealed it was behind the project all along.

It plans to open two 450,000 sq ft (41,800 sq m) data center halls, as well as a 70,000 sq ft (6,500 sq m) administrative building. Construction is expected to take about 18 months.

Facebook's Furlong said that the data center campus has been in the works for years, while Governor Ricketts added that it will create hundreds of (temporary) jobs. An earlier proposal put the number of construction jobs at 150.

The Sarpy data center is the second Facebook has announced this year, coming a few months after it confirmed it plans to build a facility in Odense, Denmark.

http://bit.ly/2oy9qrm

Contractor reveals Apple's $50m Nevada data center plans

Apple is planning a $50 million data center codenamed Project Isabel at the Reno Technology Park in Nevada, where it already has one data center, and another in construction, according to reports.

A contractor filed an application to build the new 373,000 sq ft (35,000 sq m) facility earlier this month - and then withdrew the application on the same day, according to *Fortune*. Apple first built Project Mills on the Reno site, in 2012, and has applied to build a new data center codenamed Project Huckleberry there.

The contractor then promptly withdrew the application, but it can be resubmitted. No other details are available.

http://bit.ly/2njQ9Wg

Switch opens the Pyramid campus in Michigan

American cloud and colocation provider Switch has opened the first phase of what's expected to become the largest data center campus in the Eastern United States, located in Gaines Township near Grand Rapids, Michigan.

Switch Grand Rapids facility was designed to Tier IV Gold standards and is powered by electricity from 100 percent renewable sources.

The first phase of the project has delivered more than 225,000 square feet (20,903 sq m) of white space, located inside the pyramid building that was previously housing the headquarters of Steelcase, at one point the largest office furniture manufacturer in the world.

The campus is set for rapid growth, and will eventually offer 1.8 million square feet of white space and up to 320MW of power capacity, at a cost of approximately $5 billion.

http://bit.ly/2obvkR8

The Kingdom of Bhutan opens first government data center

The Kingdom of Bhutan has opened a 2,500 square foot (232 sq m) data center at the Bhutan Innovation and Technology Centre in the Thimphu TechPark.

The facility, designed to Tier II standards, already has 22 government services running on it, which take up 60 percent of the current storage capacity of 50 terabytes.

The Nu120 million ($1.8m) facility was funded by India, under its Project Tied Assistance initiative. For the latest PTA, India has committed Rs28,000,000,000 ($430m) in assistance, funding 83 projects, including the e-Governance Program and national broadband master plan that this is a part of.

Tied aid is the controversial practice of giving financial aid that must be primarily spent on goods and services from the country giving the aid.

http://bit.ly/2n9QZES

Google announces three new cloud regions: California, Montreal and the Netherlands

Google confirmed plans to launch three new cloud regions - in California, Montreal and the Netherlands. It currently operates six regions, but plans to have more than 17 locations 'in the future.'

Each region will have a minimum of three zones - geographically diverse data center locations. The search giant did not give a timeline for the rollout of the regions, but it did open a €600 million ($635m) data center in the Netherlands last year.

This year it plans to launch new regions in Mumbai, Singapore, Sydney, Sao Paulo, London, Frankfurt, Finland and North Virginia.

Brian Stevens, VP of cloud platforms, said: "These new regions will deliver lower latency for customers in adjacent geographic areas, increased scalability and more disaster recovery options."

http://bit.ly/2mJp8Ox

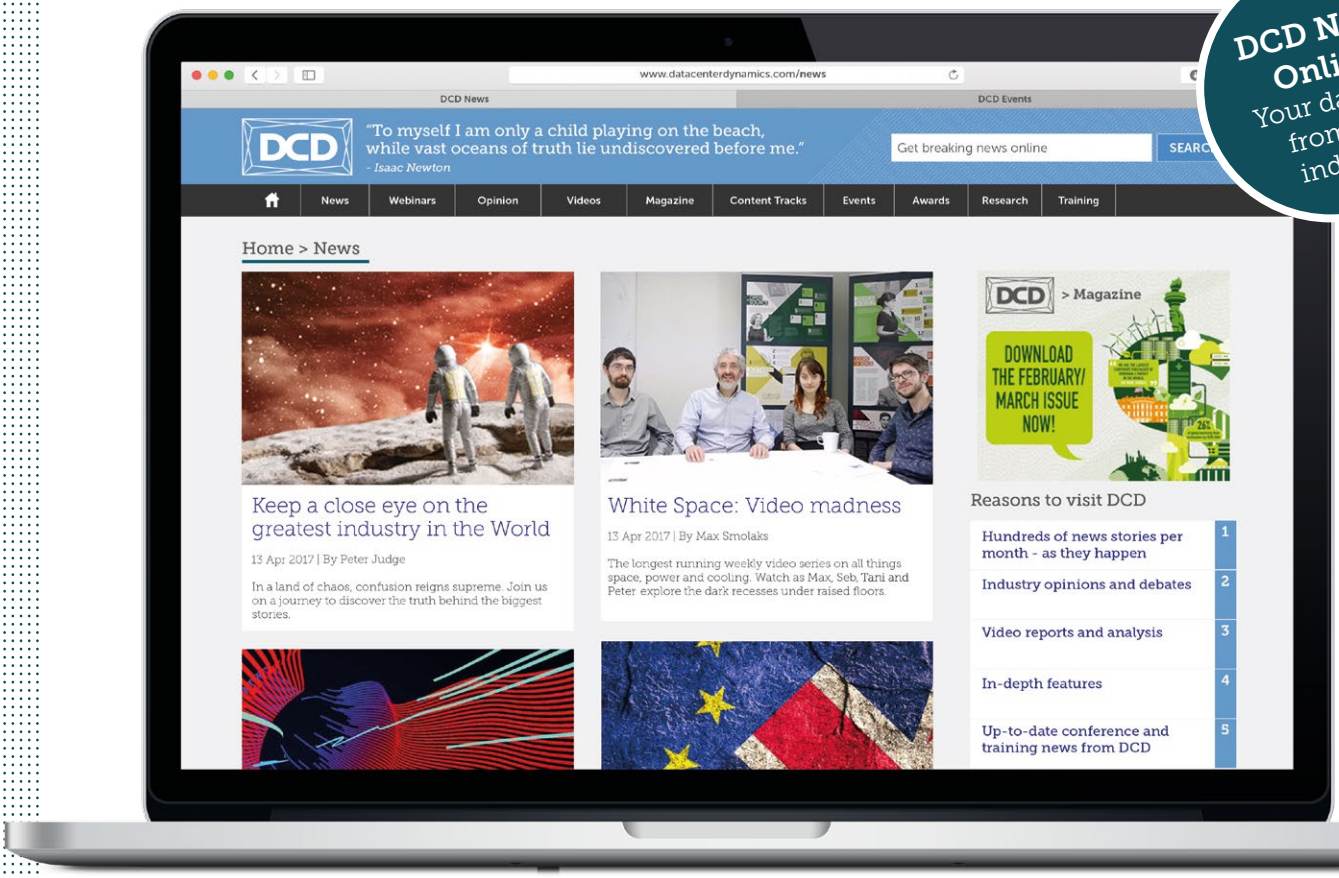

DCD News Online
Your daily fix from the industry

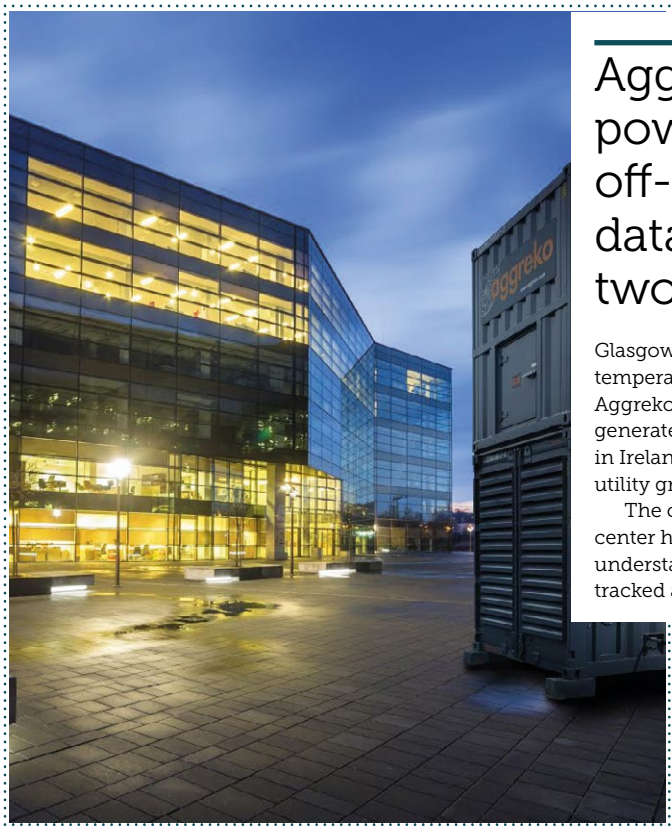

Aggreko to power mystery off-grid Irish data center for two years

Glasgow-based power and temperature control company Aggreko will supply 14MW of gas-generated power to a new data center in Ireland until it moves to the main utility grid network in 2019.

The company behind the data center has not been named, but *DCD* understands that the client has fast tracked a project, causing it to need power ahead of schedule.

Aggreko will also supply an additional 4MW as contingency power for when repair and maintenance is undertaken, for a total of 18MW.

"Data centers are being constructed at such a speed that in some countries the local infrastructure just cannot keep up with demand," Billy Durie, head of European sector and account development for the company, said.

"Temporary gas-generated power makes perfect sense for data centers that need to be operational before a connection is available from the local power supplier, or simply where there is not enough capacity from the grid."

http://bit.ly/2nwbBs1

Peter's random factoid

Facebook currently consumes around 7.5 quadrillion instructions per second. That's one million instructions every second for every single person alive.

Chinese CDN ChinaCache to sell its data center business

Chinese provider of content delivery network (CDN) and cloud services ChinaCache is to sell a 79 percent stake in its data center business, ChinaCache Xin Run Technology, for RMB221.2 million ($32.1m) in cash before fees and expenses.

Xin Run will be bought by Tianjin Shuishan Technology, Shanghai Qiaoyong Equity Investment Fund Management Co., and Tianjin Dingsheng Zhida Technology Co, with the companies acquiring 47.67 percent, 26.33 percent and 5 percent of Xin Run, respectively.

"We believe the transaction will enable Xin Run to achieve growth and unlock value, which may further benefit our shareholders. After the transaction, ChinaCache and Xin Run may explore partnership opportunities to provide our enterprise customers with premium total solutions," Song Wang, Chairman and CEO of ChinaCache, said.

In addition to being CEO of ChinaCache, Wang controls both Tianjin Shuishan and Tianjin Dingsheng.

Tianjin Shuishan, which will become the largest owner of Xin Run, will loan money from Shanghai Qiaoyong or its affiliates to finance its part of the acquisition.

http://bit.ly/2oyU9T0

BASF picks HPE to build 1 Petaflop chemical research supercomputer

The world's largest chemical producer, BASF, will work with Hewlett Packard Enterprise to build a supercomputer for industrial chemical research at the company's Ludwigshafen headquarters.

BASF previously turned to HPE for its data center needs, in 2015 outsourcing its two Ludwigshafen data centers to the US IT company, and transferring roughly 100 jobs.

"The new supercomputer will promote the application and development of complex modeling and simulation approaches, opening up completely new avenues for our research at BASF," Dr. Martin Brudermueller, BASF CTO, said.

The supercomputer will be based on the latest generation of HPE Apollo 6000 systems, and feature Intel Xeon processors, Intel Omni-Path Fabric and HPE management software, with an effective performance of more than 1 Petaflop.

http://bit.ly/2mQf0EZ

Ford to build $200m Michigan data center

Ford plans to build a $200 million data center in Flat Rock, Michigan, its second new facility after its previously revealed Dearborn, Michigan data center.

Details on the size and contents of the data center are lacking, but Ford said that it expects its data storage requirements to increase from its current 13 petabytes to more than 200 petabytes in 2021.

This expected jump in storage requirements comes as Ford tries to rebrand itself as a mobility company, as its core automobile business comes under threat from Silicon Valley's self driving endeavors and the rise of ride sharing.

In 2016, Ford and Baidu invested $150 million in Velodyne, a manufacturer of lidar (light-based ranging) sensors used in self driving cars. The company then announced that it planned to have a fully operational self driving car - without a steering wheel - by 2021, although it is likely the car would be used only in certain city centers and see its speed limited.

Software for the car will be developed in part by Argon AI, an artificial intelligence company founded by the ex-heads of Google's and Uber's self driving divisions. Last month, Ford announced that it would invest $1 billion over the next five years in Argon, becoming majority shareholder.

"There's a war for talent out there," Ford CEO Mark Fields said at the time.

http://bit.ly/2nynyxp

Facebook refreshes its servers at OCP Summit

Facebook has announced a complete refresh of all of its servers, and shared details of the new range with the Open Compute Project, the open source hardware group which it established in 2011.

The four models include a storage server, two compute servers and a specialist appliance designed to train neural networks.

The Bryce Canyon storage server has a 20 percent higher HDD density and a fourfold increase in compute capability over its predecessor, Honey Badger.

The Yosemite v2 compute server is intended for scale-out data centers. It's designed so that the sled can be pulled out of the chassis for components to be serviced, without having to power the servers down.

Tioga Pass is a new compute server with dual-socket motherboards and more I/O bandwidth for its flash, network cards, and GPUs than its predecessor Leopard.

Big Basin is the server that trains neural networks - it replaces Big Sur, offering more memory and processor power.

Vijay Rao, Facebook's director of technology strategy, said: "Open hardware increases the pace of automation. It makes it possible for everyone to work at the speed of software."

http://bit.ly/2mbCwrw

Microsoft shows ARM data center servers

Microsoft has announced support for ARM based servers in the data center, including them in the latest stage of its Olympus Project open servers, and even porting Windows Azure to run on ARM.

ARM-based chips from Qualcomm and Cavium are in two of Microsoft's new Olympus "universal motherboard" designs, available as open source, Microsoft's general manager for Azure hardware Kushagra Vaid said. The company also ported Windows Azure to run on ARM - but only for internal use, rather than for external customers.

"Microsoft has accelerated the timetable for cloud data center ARM processor adoption," commented Paul Teich, principal analyst at Tirias Research to *DCD*.

"We are still six months before Cavium and Qualcomm deliver these chips to market, and this will get other cloud giants interested. It completely scraps the IDC and Gartner forecasts for ARM servers - and it's a good move for Microsoft."

Microsoft's Project Olympus, its open source blueprint for hyperscale server hardware, was first announced at DCD Zettastructure last year.

http://bit.ly/2okfH6J

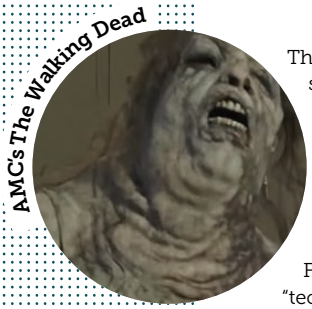

AMC's The Walking Dead

5,600

The number of US federal data centers still remaining after seven years of consolidation, down from over 10,000 (Government Accountability Office)

http://bit.ly/2nmVfRB

Nielsen's TV audience ratings delayed by data center outage

The data center of broadcast analytics specialist Nielsen suffered an outage in mid-March, causing ongoing delays in TV ratings for programs aired on a number of networks over a weekend.

The facility, located at the company's Global Technology and Innovation Center on 1 Nielsen Way, Oldsmar, Florida, initially reported an unspecified "technical issue" at 9am ET on Sunday. The company then told clients it had indeed suffered a blackout at 11am ET, but the data center power was back on.

Data was collected despite the outage, but the systems had to be rebooted in order to process and generate the ratings, causing the delays. Nielsen is yet to explain how a power issue could disrupt its servers, which should have had backup power in place.

Ratings trickled out days later, but the delivery remained slow and intermittent. Although delayed TV ratings aren't particularly dramatic for network ad deals, which are negotiated based on data collected up to a week after airing, the incident does not bode well for the company, which has been criticized for failing to adopt to modern viewing habits.

According to *The New York Times*, $70 billion in advertising dollars is traded in the United States every year solely based on Nielsen's ratings, but networks appear willing to leave the company behind if it fails to step up its game.

http://bit.ly/2obzh8a

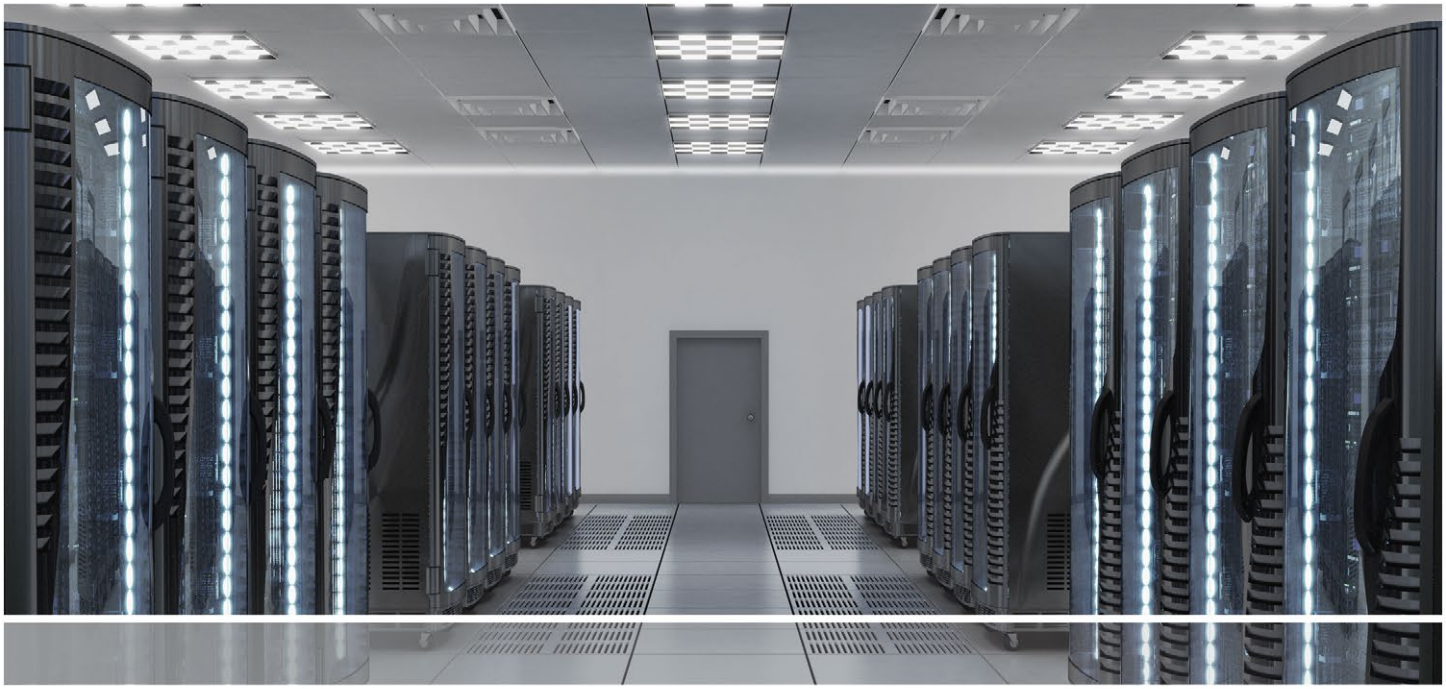

SAFE & RELIABLE DATA CENTERS

Connect with Us

Safe and reliable connector solutions from Stäubli Electrical Connectors.
With an increasing trend in the number and size of data centers, the call for energy efficiency gets louder. That means data center components need to be more energy efficient and increase use of power sources with a small carbon footprint such as solar or wind power. Equipment needs to be safe, reliable and designed for quick replacement.

Are you looking for safe and reliable connector solutions that will help to lower down times of your data center and make it more energy efficient?

We offer innovative connector solutions featuring low contact resistance and very little power loss:

Powerline
Insulated single power round connectors rated up to 1000 V / 1000 A with various termination options for power distribution.

Slide-inline
Fork connectors for busbar applications in server cabinets or overhead bus ways.

CombiTac
Modular connector with multiple options including power, signal, data and fluid connections. Ideal for rack systems and applications requiring liquid cooling.

www.staubli.com/electrical

Powerline

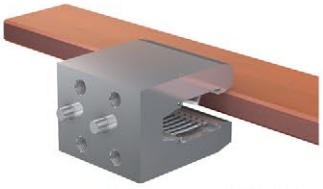

Slide-inline

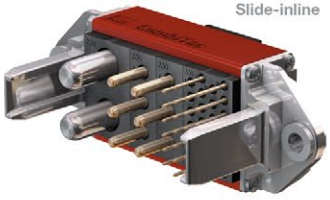

CombiTac

Multi-Contact

MC

STÄUBLI

Stäubli Electrical Connectors, Inc. U.S. Headquarters 100 Market Street Windsor, CA 95492/United States
Phone: +1 707 838 0530 ec.us@staubli.com

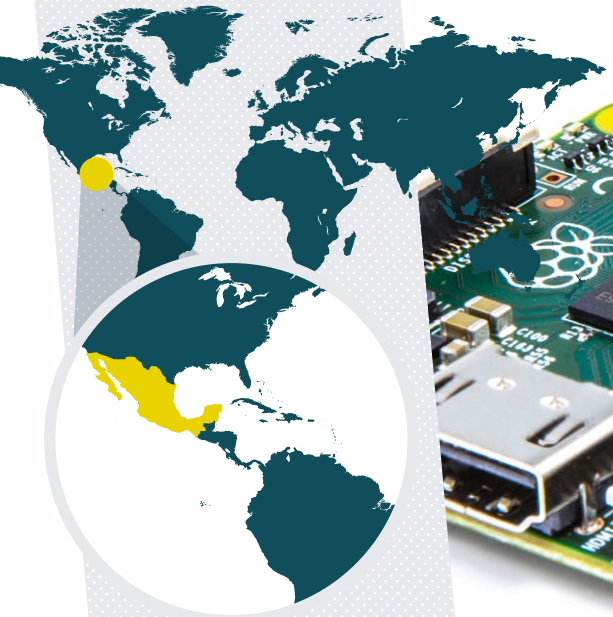

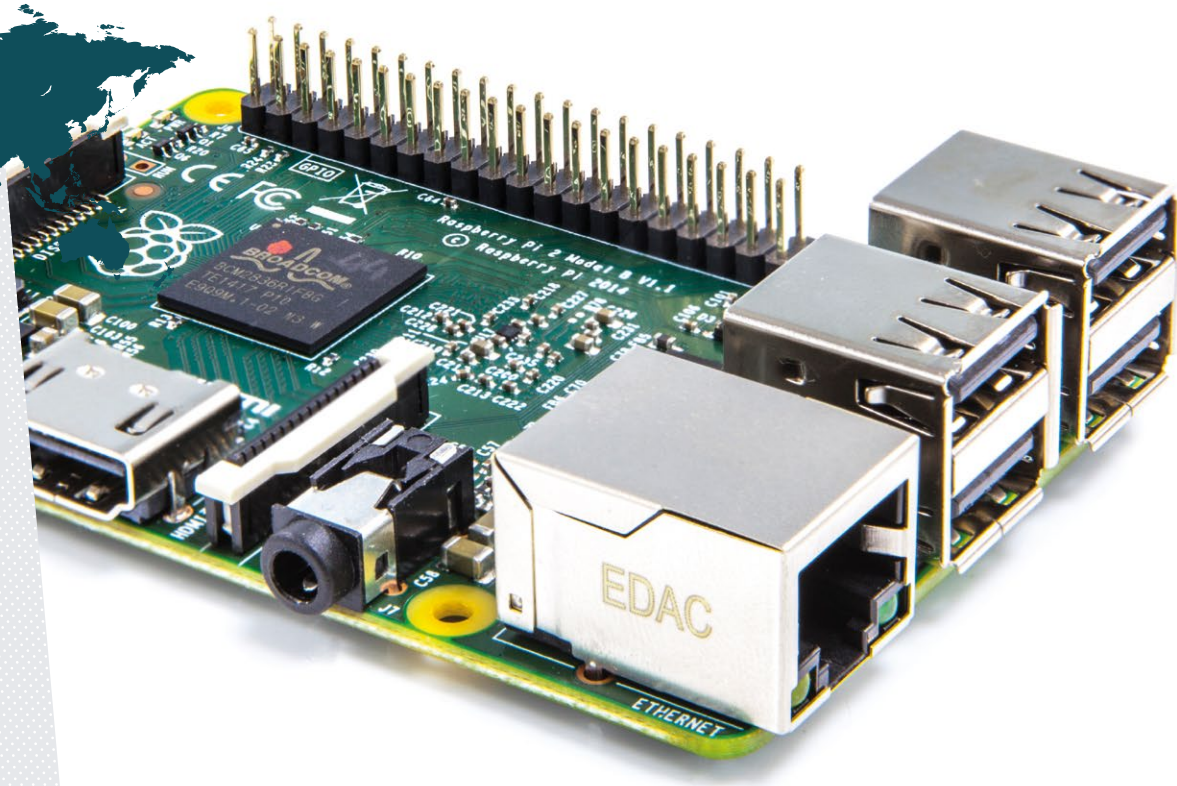

The project at a glance

Team

- Rogelio Garcia Cabrera, master of computer science
- Arturo Benitez Mejia, master of innovation and renewable energy
- Dzib Jesus Jose Sanchez, mechatronics engineering
- Diego Sabino Hernandez, electronic systems engineer

Development period

- January 2015 to July 2016

Load monitored

- 650kW

Data center

- Facility Grupo Salinas in Mexico City

Hardware used

- Arduino and Raspberry Pi

Advantages

- Modularity, sustainability, low cost, high speed of deployment, no licence required

Open source IoT protects Mexican bank

Grupo Elektra-Banco Azteca is using Raspberry Pi and Arduino to predict and avoid equipment failures. *Virginia Toledo* finds out more

Virginia Toledo
Editor LATAM

Banco Azteca is one of Mexico's largest banks, with operations in six Latin American countries. It was created in 2001 as a financial services subsidiary for the retail firm Grupo Elektra, which was set up in 1950 by Hugo Salinas Rocha.

To reduce the risks of failure in the bank's data centers, Grupo Elektra-Banco Azteca's vice president of systems Manuel Gonzalez came up with the idea of a monitoring system which could detect dangers and warn of possible failures. The bank built a detection system in-house, and launched it first in the Grupo

Salinas data center in Mexico City in early 2015. This project was completed in July 2016, and the bank plans to roll it out to more data centers.

The system is intended to assess the quality of energy delivered, the quantity of energy consumed and other metrics such as temperatures across the site.

"One of our main requirements was that these new systems were not independent, but quite the opposite. They must be able to be integrated into a single solution, so that in the future, one system can give us information about as many factors as possible," says Rogelio Garcia Cabrera, director of data centers in Latin America for Banco Azteca.

The bank analyzed several market-leading data center infrastructure management (DCIM) products but none of them met its specific requirements, so the company decided to create its own tool. The result is not a traditional DCIM app, but one tailored to the company's needs.

The project's main innovation was to use open source software, and cheap high performance system-on-a-chip (SoC) microcontrollers and microcomputers - specifically the Arduino and Raspberry Pi.

Open source is becoming more popular in the industry since it allows for constant improvement, while knowledge and new developments are shared with anyone who wants to use them.

Garcia Cabrera's team developed modules which could be connected to the mission-critical equipment, in order to provide an early event detection system.

The system developed by the team monitors energy delivered in real-time, determines the energy consumption of hardware, and gives an indication of the amount of wear in the cooling systems. It also offers a look at temperature levels and humidity inside the site.

"The advantage of a homogeneous system like this is that it allows us to determine how long a condenser has been overheating," says Garcia Cabrera. The condenser may be overheating for some time before it causes a temperature rise in an area within the facility. The useful life of the condensers will be affected by the length of time they are running too hot.

Implementation of the system is quick. In just a few minutes, sensors to measure current, voltage, temperature and humidity can be installed, connected to processor cards, and start transmitting information to a database and a real-time display panel.

It is also easy to remove and repair sensors and cards, because of their low cost. The cards and sensors are easy to buy in electronics stores and, if necessary, the group could make its own cards and be independent of suppliers. This is all due to the work of the open source community which publishes the schematics for free, in an official repository.

The arrangement of sensors and circuit boards is not static, as it allows the group to integrate more advanced cards without modifying the sensors. For example, information can be transmitted using GPRS, but by simply changing the communications module, it can be upgraded to 3G technology.

Data center managers do not need to be passive consumers of products anymore

This structure also makes the system more secure than commercial systems on the market, its creators say. There has been a surge of attacks based on unsafe Internet of Things (IoT) devices, but the bank's early warning and monitoring system can be kept safe because it is possible to update the firmware remotely. As a new security breach is identified, it can be solved and corrected long before all the devices are compromised.

"This information gives us a new opportunity to not only detect early, but to predict the behavior of the infrastructure," Garcia Cabrera says.

At the moment the system is being introduced to more and more equipment in the Banco Azteca's facility, and the information gathered will allow the team to identify patterns of behavior. This data, combined with prediction algorithms, can provide an accurate model that can foresee equipment failures before they happen.

This will lower maintenance costs and help avoid unscheduled downtime. Eventually, the system could be fine-tuned using the tools of advanced analytics, big data and machine learning.

With this kind of development, data center managers do not need to be passive consumers of products anymore, but can start to develop technology according to their needs, says Garcia Cabrera: "This will be considered one of the most important applications of the IoT. This change is taking place in the most advanced data centers around the world, and a project like this will bring on the so-called Fourth Industrial Revolution." ●

Indonesia ups its data center game

An upgrade to the power infrastructure and new financial data regulations could be just what the country needs, says *Paul Mah*

Paul Mah
SEA Correspondent

With an estimated population of 263 million, Indonesia is the fourth most populous country on Earth and the largest economy in Southeast Asia. Data centers there have some ground to make up, but they are expanding rapidly.

Facilities built in Indonesia now are more reliable, and operators are working hard to meet booming demand while complying with regulations, although sourcing reliable power and trained staff remain key problems, according to Alvin Siagian, vice president and director of NTT Indonesia.

"Data centers built in this country prior to 2009 are Tier I and Tier II. However, companies are starting to understand the benefit of data centers. Businesses that were previously immune to frequent downtimes are starting to realize they can no longer avoid or pretend to avoid the importance of data centers," said Siagian.

This is due to developments such as e-commerce, on top of more frequent Internet use, and upcoming trends such as the Internet of Things (IoT) driving demand for reliable colocation.

In 2014, some 15 percent of the first two floors of NTT's Jakarta 2 data center was in use. Now, the fourth floor of the eight story building is currently being filled up.

Another consideration is stricter requirements for the financial sector, thanks to the implementation of Government Regulation 82 of 2012 (PP82/2012), which prohibits financial data from being kept outside the country without prior approval.

The regulation calls for each vertical industry to adopt its own set of rules. In November 2016 the Financial Services Authority of Indonesia introduced its own sub regulation called IT Risk Management, that adopts PP82/2012 in its entirety. This will increase the demand for services from local data center operators.

"Data centers [serving the financial sector] must also abide by banking rules and must meet IT compliance rules. A lot of providers are not prepared," he concluded.

Indonesia's power supply is unstable, with rolling blackouts caused by shortfalls. Uneven electricity distribution was seen as the main challenge in Indonesia by an earlier IDC data center index of the Asia Pacific region.

With a country-wide undersupply and an electrification ratio of just 74.4 percent, the energy situation is compounded by Indonesia's complex geography with thousands of islands, and increased consumption that has surged by as much as 50 percent in the decade up till 2013.

The government is working to reduce energy and fuel subsidies to free up more of the budget for infrastructure - and now has an ambitious expansion plan to build 291 generation plants, over a thousand substations, and 47,000km of new transmission and distribution lines, according to a report by the Lowy Institute for International Policy.

Siagian hopes power privatization could ease matters. At present, state-owned utility Persero-Perusahaan Listrik Negara (PLN) remains the sole purchaser of power output and has a monopoly on the transmission system.

In the meantime, providers have options in the form of independent power stations, such as DCI Indonesia's JK1 data center which has access to the power generating plant in Cibitung Industrial Estate.

> *On paper a self-built data center in Indonesia can look cheap, but costs are deceptive*

Costs in Indonesia can be deceptive, says Siagian. Thanks to the comparatively low cost of land and building in Indonesia, it can look a lot cheaper on paper to build a self-operated data center than to go down the outsourced route.

In fact, merely tallying up the land and building cost does not offer a realistic representation of the TCO (total cost of ownership), said Siagian. This is because such an approach leaves out the cost of ensuring uninterrupted power, the cost of hiring the skilled staff, and the business cost of outages.

The importance of hiring professionals with the right operational experience cannot be overemphasized. Drawing upon his decades of experience in the IT field, Siagian noted: "Data centers used to revolve around M&E (mechanical and electrical infrastructure) in the past. But running a critical data center today, we need to understand the IT side of it. Where do you put your server, and how do you scale it up [or] scale it down?"

There is also room for outsourced data center operators to improve, too. "Some of your applications, or lines of business, don't need a high level of uptime. You need to know how you benefit your client so that you can scale up and down," he said. Siagian also pointed to private cloud infrastructure such as hyper-convergence with its power dense hardware, and said: "It is about scaling up, and yet being able to maintain the uptime."

It is clear that Siagian is passionate about hard deliverables such as uptime. NTT Indonesia is a member of the Association of Data Center Providers (IDPRO), which is currently working with the government to help design a framework for standards and service level agreements (SLAs) that can be widely used.

"I [would] rather work by offering an SLA, and standing by it. Tell me what your penalty is, and let me prove it to you," said Siagian.

Indonesia

Capital

- Jakarta

Population

- 263 million

GDP

- $940 billion

Energy

- Indonesia exports coal, gas and oil. It has good resources in renewables, including wind, hydro and geothermal power.

The lights are going out in data centers - so why is there anyone still at home?

POWERED BY PEOPLE

Data centers are now so fully automated that they run almost unattended. They don't need people anymore, and their basic environmental needs seem to have diverged from those of human beings. Do data centers still need people in them?

For day-to-day operations, IT hardware is better off in an environment that isn't designed around people (see On life support). It's also cheaper to build facilities out of town, and close to sources of power. But data center operators say they still need warm bodies.

"No matter how much investment you make in construction, if you didn't invest in addressing the human part of the operation, you could devalue the whole investment from day one," warns Paul Saville-King, managing director of CBRE Data Center Solutions, a division of the specialized real estate firm CBRE.

620k people work in data centers (DCDi, 2015)

Data centers designed for reliability will only actually operate at a high reliability if they are staffed by people who are able to handle the demands, says Saville-King: "You can have a facility that certified for its Tier IV design, but it's at risk based on the people - unless you invest in training and preparedness and responsiveness."

One simple factor to encourage good people is to be where they want to be. Data centers work better when they are close to people, says Andrew Fray, UK managing director of colocation provider Interxion: "The only way I can explain it is people exert a gravitational pull. In this world of amazing technology, it's still important to have human contact."

Fray's urban data centers couldn't be closer to people. On London's Brick Lane, the key amenities are easy travel, excellent curry restaurants, and 24 hour bagels in the same street. It seems food is a factor attracting business from the comparative sterility of nearby Docklands.

Even Slough, the out-of-town hub which was decried as a soulless wasteland by John Betjeman and celebrated by David Brent, has one big draw for humans: its proximity to Heathrow Airport.

Peter Judge
Global Editor

Advocates of the cloud dismiss these colocation facilities and their customers as "server huggers," people who are fixated on their IT at the expense of their real business. But there are industries that still see benefit in installing and configuring their own equipment - and some organizations where regulations or policies actually demand it.

IT is still changing fast enough that hardware upgrades are sometimes required. And systems running remotely will still have to be audited. This will mean actual physical checks on what is in the racks.

So data centers still have to be designed for human beings to access. A large webscale facility may be virtually automatic, with batch-driven hardware upgrades, but there's still a handful of people there.

Colocation spaces have more bustle. The space provider, and any networks and cloud operators will need to come in, while end users will need to tend the IT equipment in their racks.

With people coming and going every day, a colo campus needs complex security - so that each set of users can always access the kit they need to get to, but everyone's equipment remains private. Biometric access is normal at all levels from the site down to corridors halls, cages and individual racks.

A popular colo will host a virtual community, with business partners making fast links between servers located in the same site (so-called "East-West" traffic). Meeting rooms and office facilities can make this physical, giving executives and IT staff a place to plan their collaborations.

There are issues to deal with, however. Data centers are designed around the ▶

On life support

IT systems don't want the same things as people. They generate heat, and it's far more efficient to let that heat warm the aisles up, rather than waste energy cooling it down. Downstream of a rack of air-cooled servers, the temperature should be distinctly uncomfortable for a human being.

Servers don't need light to see by, or kitchens, and the only plumbing they need is for any water cooling the more advanced systems might be using. As cloud systems have become more automatic, some data centers have moved to lights-out operation, shedding much of the life support humans demand.

Electronics are better off without oxygen, which makes fire and corrosion possible. In an extreme example, Microsoft's Project Natick has tested a micro data center operating underwater. The rack is in a pod that's too small for a person to get into, even if it had an airlock, and it's full of nitrogen.

On the sea bed, Natick operated without human intervention for months on end. Any servicing and repairs had to wait until it was hauled up to the surface.

"No matter how much investment you make in construction, if you don't address the human part of the operation, you could devalue the whole investment from day one," Paul Saville-King, CBRE

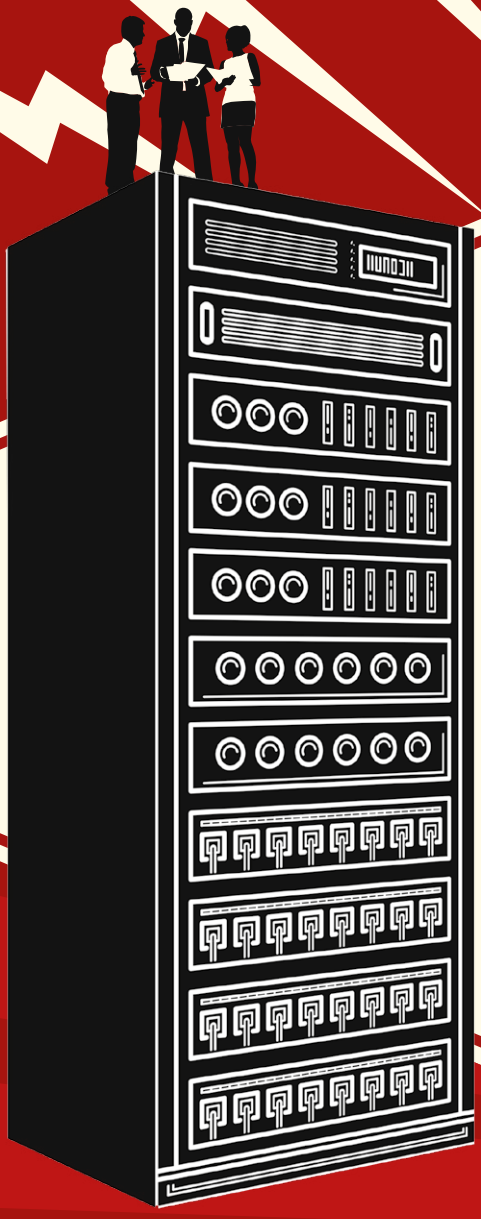

needs of servers, so anyone entering that space will have to watch their step. Raised floors and equipment can be a hazard, as well as the presence of high voltage circuits.

There should be induction courses. You may be fully aware of the risks, but if a data center lets you in without a reminder of the ground rules, be wary. A health and safety violation by a less-experienced visitor or staffer could lead to an outage for your services.

Despite this, human factors programs are still the exception, warns Saville-King. Training should not just be on the technologies, but also cover "soft skills," he says. These include communications skills and assertiveness, both vital when a problem develops and cooperation is needed to deal with it. "We must develop people that can make good decisions under pressure," he says. "They must be properly trained, repeatedly drilled, and given practice at working under varying conditions."

The overall job is changing. The sector is still expanding rapidly, but the number of people required to operate a data center is going down, with thirty people or less running a huge webscale facility.

"A lot of people build and commission a data center. Surprisingly few people work inside it," says Peter Hannaford of specialist staff consultant Datacenter People. Data center companies need C-level executives and sales people, but recruitment there goes in fits and starts, he says. Around 60 percent of jobs are for operational engineering staff.

Thanks to the continued growth of data centers, there's high demand, and staff are being drafted in from other critical environments, such as oil and gas and the military. But while demand is steady, the required skills are changing, observes Hannaford: "There's a new breed of data center technician, or data center engineer, with multiple skills."

2015
the year when IT outnumbered facilities (DCDi)

There are two major problems with recruiting people for data centers, according to Dean Nelson of Uber Compute. Firstly, the people who would be great for the job aren't aware of data centers, and aren't applying. And secondly the skills are not fixed. If you hire someone with deep skills suitable for today's technology, they will become much less useful when that technology becomes obsolete, which could happen at any time.

"Most people don't know this industry exists," says Nelson, "and we don't have the university exposure, so people aren't excited about jobs in this field." Instead of data centers, students are going into other jobs in business, computer science, or programming.

Nelson's answer to the need for human skill is to professionalize. He set up Infrastructure Masons (see Are you an

"Most people don't know this industry exists, and we don't have the university exposure, so people aren't excited about jobs in this field," Dean Nelson, Infrastructure Masons / Uber Compute

Infrastructure Mason?), a professional body which aims to encourage everyone in the data center sector to up their game, with an awareness of the responsibility they have for the world's data and its infrastructure.

The people they are looking for may not be in data centers right now, says Nelson: "We've been looking for people with experience in different industries, and bringing them in to look at our problems."

Sometimes people from adjacent industries bring invaluable experience, as Christian Belady, head of cloud infrastructure strategy at Microsoft, found when he hired a senior executive from the dairy industry. Unlikely as it may sound, her experience handling a supply chain for a perishable - but valuable - commodity was directly applicable in data centers, and her experience in an environmental plant also came in useful.

Mechanical and electrical systems used to have their own staff. Now the roles are combined in one person, and the work is becoming more IT related, for two reasons. "It is easier to learn about mechanical and electrical systems than it is to learn the IT bit," says Hannaford. And the physical infrastructure is now monitored and controlled by data center infrastructure management (DCIM), or a building management system (BMS), effectively placing the building under IT's control.

"When the service goes down, the power guy will think if there is something wrong with the power, and the mechanical guy will think about the cooling. The IT person is probably bridging the two."

If you only have one person in your network operation center (NOC) it should be a generalist with IT skills, who can identify the problem and call in appropriate mechanical or electrical experts.

"The only way I can explain it is people exert a gravitational pull. In this world of amazing technology, it's still important to have human contact," Andrew Fray, Interxion

So data center people all need IT skills now. "If you're looking at an infrastructure job because coding gives you the willies, there's bad news," says commentator Dan Rosenbaum in a blog sponsored by HPE. "You've got to get with the programming if you're going to have a career."

As well as tech, those soft skills are needed, he continues: "If you think a data center job is a good place to hide from people, the news is even worse, because soft skills are increasingly important as well."

It's such a potentially diverse role that smart firms take people with the right ▶

Are you an Infrastructure Mason?

Dean Nelson, of Uber Compute, says Infrastructure Masons is a professional organization for the people that make the digital economy work. "We build the underlying infrastructure that enables the Internet of everything. We are trying to bring this community of people together to collaborate."

The profile of a Mason is about experience, economics and stewardship. Experience relates to the Mason's years of experience in hardware, software and infrastructure. The economics factor depends on how much capital expenditure and operating expenditure a Mason has been responsible for during his or her career. The stewardship factor is about the number of industry groups the Mason has been involved with or had leadership positions at.

Collectively the founders of the group have hundreds of years of experience, and have handled many billions of dollars of capital. The group wants to encourage more people to reach this level of experience.

"This is where infrastructure professionals connect, grow and give back," says Nelson.

Join the Infrastructure Masons here http://imasons.org/join

intellect and behavior, says Hannaford: "You can change their experience and skills. You can't change anyone's intellect. Behaviors you can modify, but it's not easy. In theory you could hire someone with no experience."

Eventually a lot of this will be automated, and the number of operational staff may shrink even further. The construction staff could also decrease, as firms adopt a modular cookie-cutter approach where identical rack systems are produced in a factory-like environment.

Support and maintenance will increasingly be automated and provided by non-human systems, or by remote humans interacting through augmented reality or virtual reality, "projecting their consciousness" and operating robots where they are needed says Saville-King.

You might expect that fewer humans will increase the reliability of the system as human mistakes are programmed out of the system (see *To err is human*). But this might not be the case, warns Saville-King: "If the resilience is built into the technology, then the sort of incidents people respond to will shift to the more critical end of the scale."

One can imagine a situation where the facility runs itself, and bored operators wait for an incredibly rare instance where a decision is required. Given the importance and the rarity of those instances, how can they be trained and motivated? "People issues will be even more important in the leaner models of the future," Saville-King predicts.

None of these scenarios are the end for people in data centers, as a standardized process will stagnate. "If you keep doing what you did before, there will never be any innovation," says Hannaford. "We rely on people to innovate."

Data centers have already changed from how they were ten years ago, and more changes are afoot. Some facilities are throwing out features like raised floors and 600mm floor panels. "If a smart engineer designed a data center from scratch, it would look totally different to what we have," says Hannaford.

Humans are driving that level of change, and humans are needed to handle it. Data centers will always need creativity.

To err is human

The presence of humans implies the possibility of human error. Even the most reliable systems will eventually fail due to human error, after 150,000 to 200,000 hours.

This figure comes from the classic "*Managing Risk: The Human Element*" (2008) by nuclear scientist Romney Beecher Duffey and aviation regulator John Walton Saull. The figure can be reduced by human ingenuity applied to the system to minimize the dangers of error.

The physical design of a data center should include basic things like electrical safeguards, and clearly labelled emergency features, and the operator needs to make sure anyone in the building understands these.

A junior staff member was the last out of one facility, and found that his swipe card didn't work at the exit. Hoping it would open the door, he pressed a button. It turned out to be the emergency shutdown. He was left trapped in a dark building full of idle equipment that was losing money for customers.

Higher up the stack, human error can blight software design and operation. When the Amazon Web Services storage function, S3, failed for four hours in February, it caused a cascade of failures in web services that rely on S3. The cause was ultimately revealed as "human error:" an admin tidying up some unused virtual servers had mistyped a command and deleted vital resources, leaving S3 unable to function.

Of course, the real human error was in the design of the system that made that command possible, allowing admin tasks to have more privileges than they need - and outside of AWS, those websites and services that rely on S3 have designed in a single point of failure. That human failure cost businesses an estimated $150 million.

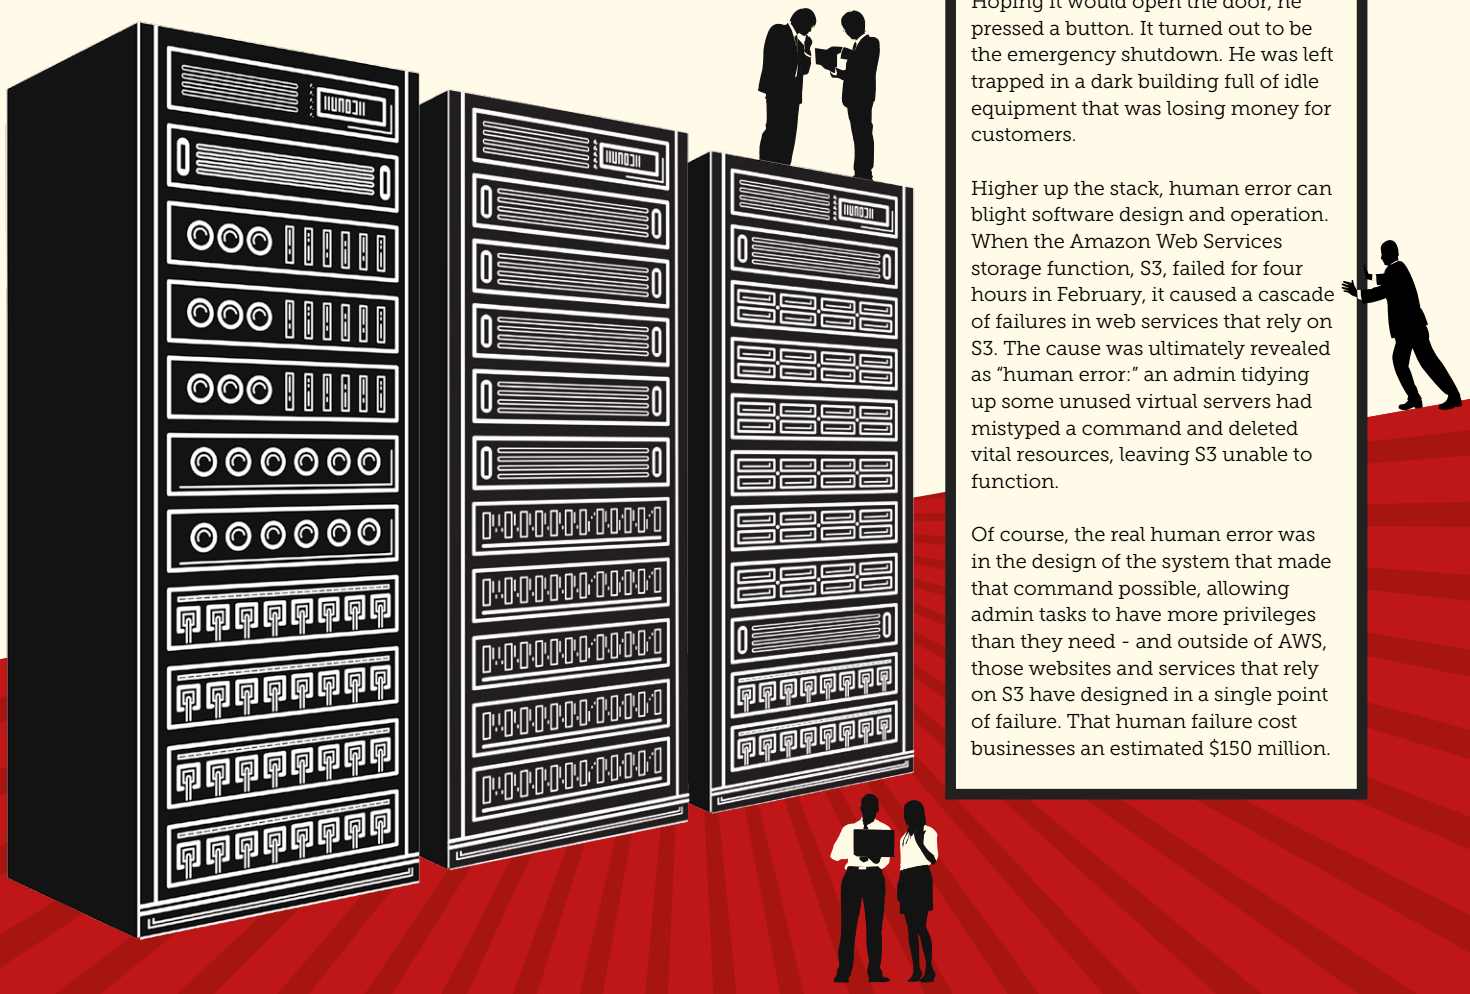

Data centers are fired by a human heart

Whatever happens to the technology, human relations are still central to innovation in the data center, says *Andrew Fray*

Andrew Fray | UK MD, Interxion

With the world mesmerized by new advances like the Internet of Things and cyber-physical systems, it is easy to forget that people, not technology, are the real driving force behind innovation.

Long before we had the cloud, robotics, or software-defined networking, rapid urbanization ignited an innovation explosion by concentrating varied industries, people and ideas into one place. 18th century London became the center for commerce and creativity because the city was a rich mix of different outlooks, skills and approaches. As the number of people, businesses and opportunities in London boomed, the city attracted even more talent and investment, creating a gravity well of exponential innovation.

In London today, I see colocation data centers as 'digital cities,' which have become a focal point for 21st century innovation. By uniting diverse digital business in a common marketplace, colocation enables innovative people to capitalize on rocketing data creation and connectivity.

Customer communities are emerging around these shared challenges, as the sheer 'gravity' of dense data, connectivity and human expertise draws companies towards the center.

With a myriad of different businesses linked together inside a single facility, colocation lets organizations develop new ideas more easily and quickly. For the price of an optical fiber cross-connect, businesses can share data, capabilities and services with a ready-formed hub of potential suppliers, partners and customers.

Colocation creates a rich marketplace platform where we can do what we all love to do: innovate.

Data centers are not simply somewhere to store your IT kit or your data; they rely on trusted human relationships to map out strategies that can flex and adapt with our rapidly changing world.

Inside the data center, access to the right people is just as important as access to the right technologies. Automation is on the rise, but skilled, clever staff remain central to a high-quality experience.

With more than 60 percent of data center operators concerned about a lack of suitably qualified people, colocating in urban areas provides access to scarce skills. Out-of-town colocation can offer upfront cost-savings, but can also sever vital relationships.

In an increasingly global world intellectual gold is still being mined in city coffee shops and restaurants, just as it was in the 18th century.

Even if the robots start to take over, innovation will still be driven by clever disruptive minds gathered at the center, not the periphery. While data centers provide us with a technical platform, it pays to remember that the data center has a human heart.

Access to the right people is just as important as access to the right technologies. Clever staff remain central to a high-quality experience

Tying down the cloud

Data center companies hope that private/hybrid interconnects will guarantee them a place in the cloud services ecosystem, says *Martin Courtney*

Martin Courtney
Freelance analyst

The slow but sure migration of enterprise applications and services into the cloud is having a significant impact on demand for data center capacity and colocation, and operators are being forced to rethink their business models.

Many have identified the interconnects which link one data center to another as a key area of growth – the chance to occupy a critical position in the cloud service chain which leaves them less likely to be bypassed in favor of direct enterprise connections into facilities owned and operated by the likes of Amazon Web Services, Microsoft, Google and IBM.

Equinix Cloud Exchange is a case in point. It is a platform which effectively replaces the direct virtual private network (VPN) links which customers might make. Instead, Equinix hosting customers get a direct link to the cloud provider using AWS Direct Connect or Microsoft Azure ExpressRoute, for example.

Cloud Exchange is touted as taking much of the complexity and management overhead out of those direct relationships, particularly where a single customer might have multiple direct connections to many different cloud service providers. It also delivers a performance advantage compared to routing traffic from the customer directly to the CSP via the Internet, said Mike Winterson, managing director of Equinix services, with an additional layer of security providing extra protection against distributed denial of service (DDoS) attacks.

Cloud Exchange does appear to be gaining some traction with customers, currently handling over 1,500 direct connections involving hundreds of thousands of physical connections, ports and switching equipment. For the moment, it handles direct connections into AWS, Microsoft Azure, IBM SoftLayer, Google Cloud, Oracle FastConnect as well as ServiceNow, WorkDay and Salesforce.com, but Equinix is planning to add other providers in the future.

Some would argue that puts Equinix firmly in the role of a cloud broker. But rather than moving towards being an aggregator itself, the company is working with other service providers and systems

integrators that want to offer Cloud Exchange capabilities under their own brand. One example is Beeks Financial Cloud, a UK-based company offering various IaaS, colocation and connectivity services to the financial sector, with Equinix estimating around 300 are deployed globally addressing additional markets such as healthcare and insurance.

"We are unlikely to move to a world where [Equinix] can address huge numbers of companies in the enterprise space, so we need to develop relationships with others who can use the lego bricks from AWS and network service providers," said Winterson.

Digital Realty launched its Service Exchange late last year in partnership with Australia-based software defined networking (SDN) interconnection provider Megaport. Like Cloud Exchange, Service Exchange is aimed at customers looking to implement large scale hybrid cloud deployments which span multiple cloud providers and locations in different countries, but need a middleman to deliver secure interconnects between public and hosted private cloud architecture, better control and more dynamic provisioning and configuration for the kind of hybrid networks they need.

Service Exchange provides direct access to the major IaaS players, including AWS, Microsoft and Google, and is working with additional SaaS players to bring their applications into the fold. Digital Realty chief technology officer Chris Sharp (who was formerly responsible for cloud strategy at Equinix) says cloud service providers themselves can use data center companies as a channel to reach more enterprise customers, saving them time on go-to-market initiatives and helping them deliver private cloud services quickly.

"We never want to look like a cloud broker, that is negatively viewed by all parties because the broker gets caught in the middle of whatever deal they have negotiated," he said. "But we spend a lot of time on the service processes to make all these clouds interact."

Data4, based in France, is also aiming squarely at enterprise demand for hybrid IT architecture, targeting early customers in financial services and manufacturing industries with high performance computing (HPC) requirements that scale out private cloud workloads to AWS' public cloud platform on demand.

"We are positioning ourselves at the convergence point between the two – a place where customers can develop and build their own hybrid IT securely and economically and have access to a portfolio of multiple [cloud-hosted applications and] services," said Data4's head of marketing Loic Bertin.

To that end, Data4 signed a partnership deal with third party cloud brokerage InterCloud in January this year. The agreement is designed to give Data4 customers direct access to hundreds of cloud service providers on a global basis, all facilitated via a single virtualized data center fabric spread across Data4's 14 hosting facilities across three sites in France, Italy and Luxembourg.

InterCloud's Cloud Delivery Platform currently connects to over twenty IaaS providers including AWS, Microsoft Azure, Google Cloud Platform, and IBM SoftLayer and in excess of thirty SaaS vendors including Salesforce.com, Office 365, Box, ServiceNow and Cisco-owned Webex. Again, Data4 does not provide the cloud resources itself but delivers the interconnects between the data center and the cloud gateway using existing points of presence (POPs).

"Of course, the CIO could buy a direct connection to each data center where the required cloud resources are hosted, but that is quite complex and he or she would have to speak to many different providers," added Bertin. "They would have to sign a dedicated contract relative to each connection, in each country."

All three companies acknowledge that continuing incursions from super scale cloud service providers are impacting their traditional colocation business, but see opportunities for peaceful coexistence and indeed mutual success through subtle repositioning.

$91.7bn

predicted size of hybrid cloud market in 2021 (MarketsandMarkets)

Why data centers want a seat at the hybrid cloud table

Analysts predict significant expansion of hybrid cloud, and/or the implementation of hybrid IT initiatives that rely heavily on different types of cloud architectures, amongst public and private sector organizations over the next few years.

In its report '*Market Trends: Cloud Adoption Trends Favor Public Cloud With a Hybrid Twist*,' Gartner outlines its belief that increased use of multiple public cloud providers combined with growth in various types of cloud services will create a multicloud environment and a need to coordinate cloud usage by utilizing hybrid architecture, for example.

Elsewhere, 451 Research ('*2017 Trends in Cloud Transformation*') predicts that CIOs will accelerate their use of AWS and other public cloud services. These will be tied into a 'blended cloud strategy' that does not lock companies into a single vendor or hosting location but matches application requirements, workloads and service requests to the best provider and data center venue.

Serious obstacles may delay those hybrid initiatives, however, not least of which are integration challenges, application incompatibilities, the absence of management tools and common application programming interfaces (APIs) and limited vendor support.

If Equinix, Digital Realty, Data4 and others offering cloud exchange or brokerage services can convince IT departments they can help overcome those problems, their own transformation from colocation and hosting providers to indispensable enterprise cloud infrastructure partners looks assured.

"We had a serious look at whether cloud was friend or foe about three to four years ago, from the board level down," said Equinix' Winterson.

In the end, he concluded that this is a major change and we are all along for the ride: "The IT industry goes through hype cycles on a regular basis – the ASP market 16 years ago was a precursor to the cloud that never took over.

"But cloud is different, you cannot stop it, you just have to transform with it or you will be sitting very uncomfortably in a few years' time."

DCD > Webinars | 2017

ON DEMAND

Making the Case for Lithium-Ion Batteries in the Data Center

Speakers:
Peter Panfil, VP Global Power, Vertiv
Tony Gaunt, Senior Director, Vertiv Asia
Tom McKinney, Director, Forsythe Data Centers
Stephen Worn, CTO, DCD

If you missed our recent webinar with Vertiv that was packed full of insight, best practices and key takeaways, then don't hesitate to watch the recording. An essential listen for data center managers globally.

▶ **Watch here:** http://bit.ly/2oDZvgt

ON DEMAND

Is Hybrid IT the new holy grail?

Speakers:
Adam Levine, CCO, Data4
Bruce Taylor, EVP, DCD

On 8 March, Adam Levine from Data4 and DCD explored the decisions needed to ensure your company is hosting your IT assets in the right place.

To make hybrid IT work effectively requires a level of computing capacity, connectivity, efficiency and resilience that most enterprise data centers simply cannot reach without prohibitive investment.

▶ **Watch here:** http://bit.ly/2n6DVEg

For more information please contact **webinars@datacenterdynamics.com**

Finding new data center frontiers

Building a data center is hard. Building a data center in Angola is harder, says *Sebastian Moss*

Sebastian Moss
Reporter

As companies around the world fight for business opportunities in established emerging markets such as India or China, others are looking further afield. But those wishing to do business in frontier markets will face issues including harsh climates, limited infrastructure, and poor levels of security.

"If you do a project in Chad, it's 50°C (122°F), you need to have the right type of guys on site," says Tomas Rahkonen, CTO of Flexenclosure, a firm which specializes in setting up data centers in tough environments.

"You need to have people who are used to solving problems under those conditions."

Flexenclosure has built its prefabricated modular data centers in countries such as Chad, Angola and the Ivory Coast. The key to pulling this off, he says, is planning, long before the equipment is shipped to its destination.

"There's thousands of details in a data center, and you can't fix stuff easily at the site. If you've missed something, forgotten something, done some bad planning, or not brought enough materials, it's difficult to go down the road and get it."

Sometimes there aren't even roads, and the modular facilities have to trek across continents to get to their destination. ▶

"With our delivery in Chad, we had four and a half meter wide modules that went through half of Africa," Rahkonen says, talking about the country's first data center, a $6m, 374 sq m (4,025 sq ft) installation.

Nick Arvanitis, global marketing director, adds: "A module on the back of a truck can get through some pretty rough terrain.

"But the challenge of getting somewhere as far away from a seaport as Chad is, isn't just infrastructural, it's also security. Particularly in sub-Saharan Africa, it's another level of challenge that one encounters in developing worlds."

Rittal also makes modules, and its director of international products, Jason Rylands, agrees: "There can be security concerns. In things I've done in places like Papua New Guinea you need to make it bullet proof just because the local people will take shots at stuff."

Transport companies and telcos often turn to security firms to protect their investment, but armed guards can do little against the elements.

"It's not that it's very hot or dry there, or it's terribly humid there," says Flexenclosure's Arvanitis. "In many countries it can actually be extremely hot and dry at times, and then extremely wet with monsoon rains and flooding, and then extremely humid."

For Chad's record heat levels, the company "specifically designed the cooling system to cope with the high daily temperatures," CEO David King says.

By contrast, the Caribbean is particularly humid, "and that can cause condensation," Rahkonen says. "You don't want hot outside air hitting cold air."

Because of this the company prefers "to do penetrations to the modules for cables from below."

He adds: "We usually have the modules on raised plinths, because there's quite commonly flooding risks as well."

And then there are the seismic issues - "we have to construct the place to survive an earthquake," Rahkonen says.

"Also insects," Rylands adds: "Insect infestations, you get some really interesting insects that like to nest in certain things."

Even if the data center is ruggedized, reinforced and ready to go, areas with unstable grids can be a challenge.

"They will need higher capacity UPS batteries, so actually lithium-ion batteries are being adopted in those sites," says Victor Cheng, VP and GM of Power System BG, Delta Electronics.

Flexenclosure's facilities in Chad, Paraguay and Angola

"You have to charge and discharge every day basically, and lithium-ion batteries are okay, if you don't completely discharge them. But lead-acid is not very fit for too many cycles of charge and discharge, so emerging markets turn to lithium."

Rahkonen says it can be that the facility "pretty much runs constantly on diesel because the grid is too bad, like in Liberia."

But before a facility can even begin to start literally burning cash away, one has to raise the funds to finance the development.

"You can have entrepreneurs that are in the countries with a business idea, but perhaps they lack a bit of cash," Rahkonen says. "So there is quite often a need to be good at working with the credit agencies and helping out with financing."

Some facilities run constantly on diesel, because the grid is too bad

Help is exactly what a lot of companies in these markets may need as "their capex is denominated in US dollars, and the local currency is going down against the US dollar everywhere, so the cost is actually increasing," Zhilei Zou, president of Huawei's carrier business group, says, referencing the surging dollar.

Despite this temporary difficulty, the market is growing rapidly, and will continue to do so, Wei Peng, VP of CBG at Huawei, believes - and in response, the Chinese networking giant set up a specialized emerging market department at the end of 2016.

But while the market may have its own unique challenges, people everywhere are the same. "As a country develops and more people become middle class," Rittal's Ryland says, "they want phones and Facebook, and it's putting a massive strain on the networks."

Last year, Kenya's 5.3 million users only needed about seven or eight racks, says Rahkonen: "It's quite small compared to developed markets, but the driver is there."

That changes things. "Suddenly it's a Western world spec, because if Microsoft wants to put their stuff in a data center in Africa, they don't want to compromise with their service quality, they want the redundancy, high ceiling heights, and whatever they require as part of their typical colo specification."

He adds: "It's actually these anchor clients that shape the spec of the data centers. That's the market logic."

In that market, telcos are leading, says Rahkonen. They seem to be the most excited, seeing a new high-growth market.

And those telecoms companies have had the luxury of watching their Western counterparts fail to capitalize on colocation. "I think the opportunity is out there for them to take this role and become the colo because their brand is so strong and it's a technology brand," Rahkonen says.

"The big names like Interxion and Equinix and Digital Realty - they're not known in these markets, but Zain is known, Tigo is known, Millicom is known, and so forth."

1 Peak experience
The Switch Pyramid, Michigan

Top 10 beautiful data centers

Data centers should look good - and this year, DCD Awards will introduce a new Design category, which honors facilities that enhance their environment. Here are ten of the best-looking data centers we know about.

If you know other sites that match these, tell us! goodlooking@datacenterdynamics.com

Peak experience

❶ *The Switch Pyramid, Michigan*

All of Switch's data centers have a trademark exterior and interior style based on patented power and cooling systems. The Pyramid site near Grand Rapids also includes an adaptive reuse of an iconic building. The seven-story steel-and-glass structure was originally created by Steelcase as a design center. Switch has opened a 225,000 sq ft (21,000 sq m) data center built into the lower two floors. With other buildings alongside, Switch will create a campus wih up to 320MW of renewable energy, 1.8 million sq ft (170,000 sq m) of data center space, and 435,000 sq ft (40,000 sq m) of disaster recovery office space.

Villain's lair

❷ *Bahnhof Pionen*

Bahnhof's data center in Stockholm pioneered the creative data center movement, when it opened in 2008 in a former nuclear bunker. Pionen's design is consciously based on a James Bond bad guy's crib, and makes several references to the 1970s movie Silent Running. Built by Jon Karlung and briefly the home of Wikileaks, it has backup power from diesel engines designed for submarines, and features waterfalls, a salt-water fishtank, and plants growing under simulated daylight.

Peak experience
The Switch Pyramid, Michigan

2 Villain's lair
Bahnhof Pionen

Sources: 1 Switch 2 Kristina Sahlén / BahnhofAB

Building altar-ations
Barcelona Supercomputing Center

Building altar-ations
❸ *Barcelona Supercomputing Center*
Opened back in 2005, the Barcelona Supercomputing Center is built in a former 19th century church. The Torre Girona was rebuilt after the Spanish Civil War and is now part of the campus of the Polytechnic University of Catalonia. Now it holds the MareNostrum supercomputer, a joint venture built by IBM and the Spanish government. For a time it was one of the world's fastest machines. It may not hold that claim anymore, but it's still one of the best looking.

3 Building altar-ations
Barcelona Supercomputing Center

Sources: ❸ *Barcelona Supercomputing Center* ❹ *datacentermurals.withgoogle.com* ❺ *datacenternavercorp.com*

DCD >Awards

Beauty is subjective, and these architectural gems might be eyesores to you. Do you like rigid functionality, decorative embellishment, or clever repurposing? As part of our campaign for better looking facilities, we are calling on the DCD Community to decide which sites look the finest.

There will be a public vote in this year's DCD Awards, for our first ever prize for data center design.

4 Google's Data Center Mural
Google, Oklahoma

5 Storage chest of data
Naver

Google's Data Center Mural
Google, Oklahoma

Google's Data Center Mural
❹ *Google, Oklahoma*
In Google's Data Center Mural project, launched in 2016, murals represent data center activity on their outside walls. In Oklahoma, artist Jenny Odell found man-made features in Google Maps' satellite images, gathering views of swimming pools, circular farms, water treatment plants and salt ponds. Working from giant cradles, 15 painters used 400 colors, transferring the images to the wall using a chalk tracing technique similar to that used by Michelangelo in painting the ceiling of the Sistine Chapel in Rome.

Storage chest of data
❺ *Naver*
Naver, Korea's leading web portal has a data center at the foot of Mount Gubong in Chuncheon, Gangwon Province, which stores its customers' online content. It's called Gak, after Kyujanggak, the royal library of the 18th century Joseon Dynasty, where Buddhist documents are stored on wood blocks. The building incorporates traditional design elements, as well as cutting-edge environmental techniques such as reuse and recycling of rainwater.

6 Under the mountain
Green Mountain, Norway

Under the mountain
❻ *Green Mountain, Norway*
A retired NATO ammunition store at Stavanger was reimagined as a 146,000 sq ft (13,600 sq m) data center by Green Mountain. The Tier III reliable underground site is inside the mountain, leaving the landscape unspoilt. The site has abundant green energy from two hydroelectric sources. Year-round cold water at 8°C (46°F) from the fjord below cools the servers, using a duplicated circulating system for reliability. The site also has a strong aesthetic. Inside, tunnels are carved from solid rock, and outside is a green and peaceful mountainscape, above the fjord.

Factory reconditioned
❼ *NGD, Wales*
NGD's Newport data center was opened in 2010 in former LG semiconductor plant, a property that had stood vacant for more than a decade. The data center space has been expanded repeatedly within the building. All the power needs are supplied by the nearby Dinorwig hydroelectric facility.

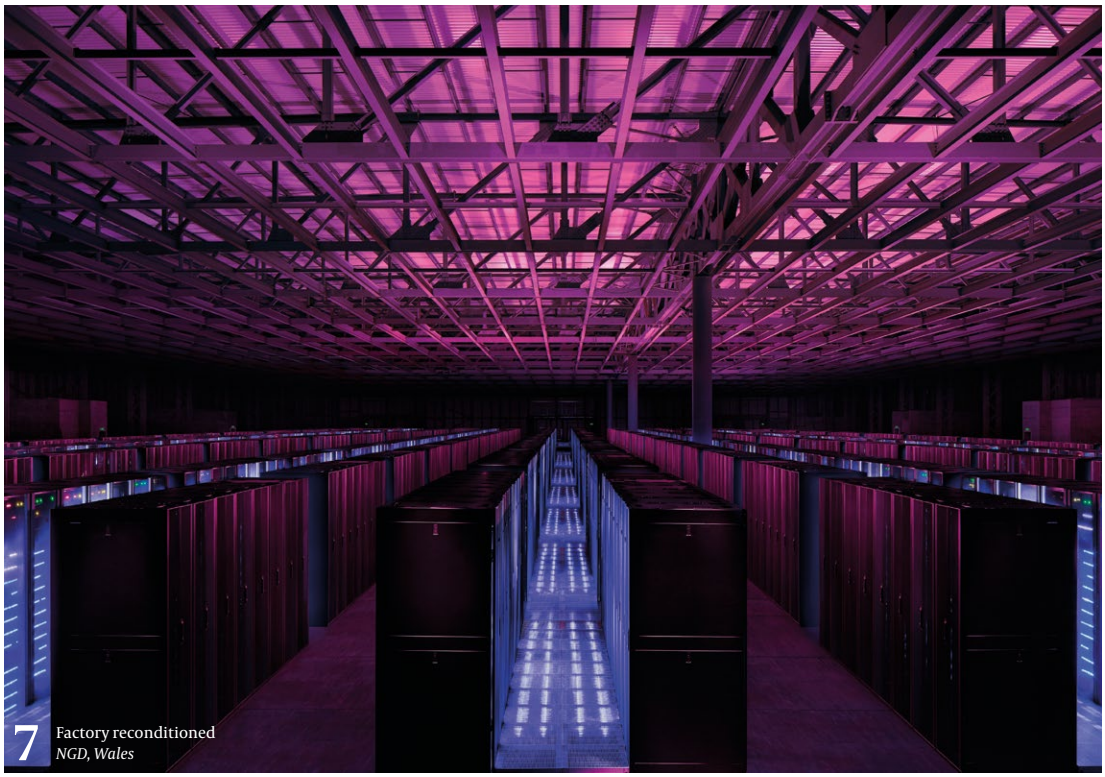

7 Factory reconditioned
NGD, Wales

Sources: ❻ greenmountain.no/Knut Bry ❼ NGD ❽ Digital Realty ❾ AQL ❿ LuxConnect

8 Towering ambition
Amsterdam Data Tower

"Data centers have a job to do, but there's no reason they shouldn't look good as well. Tell us your favorite data center so our Awards can recognize the world's best facilities!"

George Rockett,
DCD CEO and co-founder

goodlooking@datacenterdynamics.com

Towering ambition
❽ *Amsterdam Data Tower*
Digital Realty took over Telehouse's AMS 1 shortly after it was completed in 2016, opening it as the Amsterdam Data Tower. The 72m-tall building has 5,000 sq m (54,000 sq ft) of data space on 13 floors, and 9MW of power. The building uses outside air and groundwater for cooling, and also stores warm water underground. The building was designed by Rosbach Architects of the Netherlands.

A view from the gallery
❾ *Salem Chapel, Leeds*
AQL built its headquarters in Salem Chapel, the only surviving 18th century dissenting chapel in Leeds, UK. Opened in 1791, with seating for 1,000 people, the chapel closed after more than 200 years in 2001. AQL turned the ground floor of the chapel into colocation data center space with a glass roof. Above that, the chapel's balcony was refurbished as a conference auditorium. The British government launched its "Northern Powerhouse" program there in 2016, but for most people, it is better known as the place where Leeds United football team was founded.

Grass on the roof
❿ *LuxConnect, Bettembourg*
Luxembourg is a popular location for reliable data centers. LuxConnect's DC1.3 has multiple Uptime certificates giving users the option of Tier II, Tier III or Tier IV reliability. The 59,000 sq ft (5,500 sq m) building uses free cooling, and has access to all-renewable power. The building has steel mesh walls which double as a Faraday cage for security, and a grassed roof to reduce environmental impact.

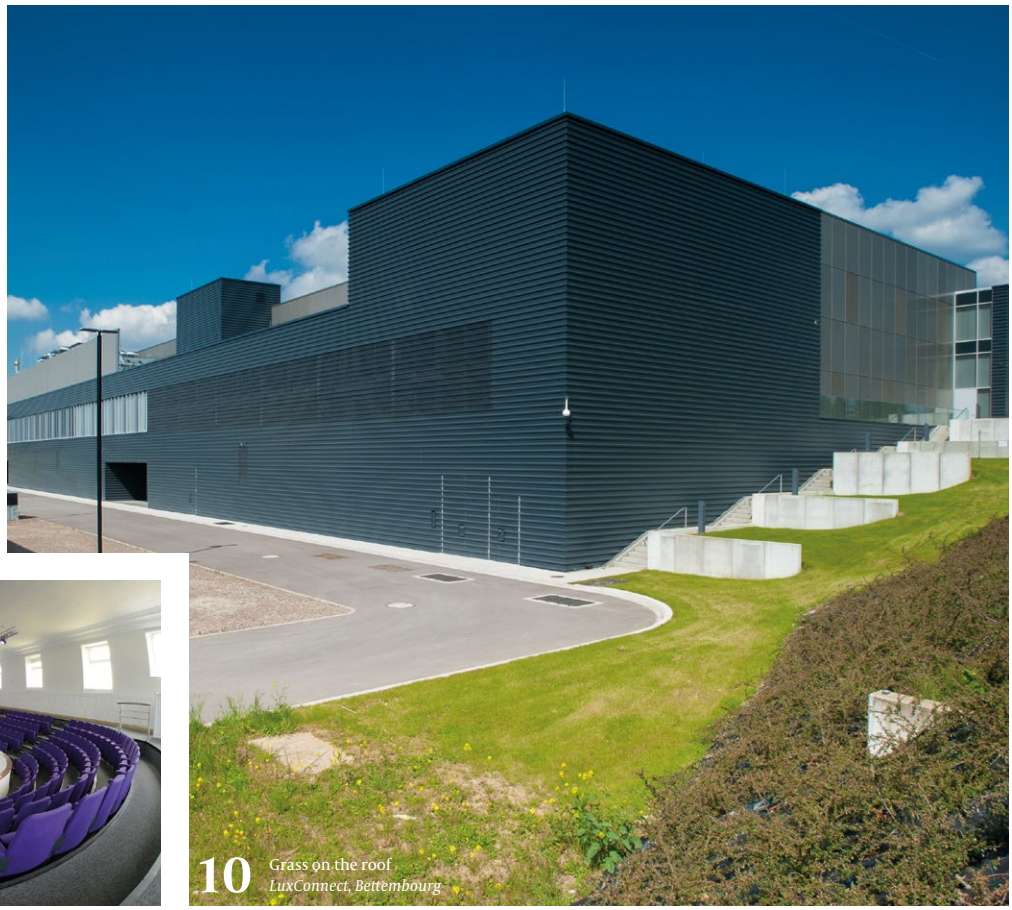

10 Grass on the roof
LuxConnect, Bettembourg

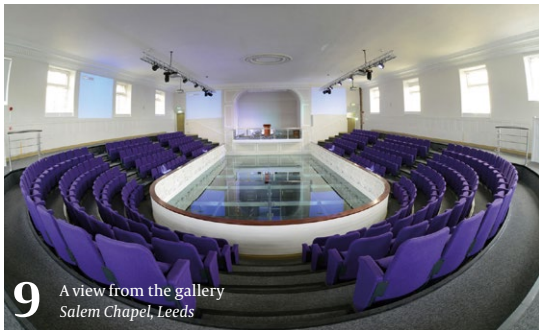

9 A view from the gallery
Salem Chapel, Leeds

The changing shape of data center markets

Data center hubs need to develop flexibility and scalability in order to remain an attractive investment proposition

Much debate has focused on 'hub' markets as major cloud, colocation and managed service players seek out new sites across the world from which to access and service customers. Hub markets have been characterized by a level of data center asset development and of investment disproportionately large in comparison to the IT requirements of the local economy. They present an economy which supports and draws from IT, and which enables them to attract regional head offices. They offer very high levels of skills, of connectivity and have an active policy of supporting and attracting foreign investment. The shift in economic weight from the established towards 'emerging' markets has positioned hub cities such as Singapore, Hong Kong and Dubai as portals into South East Asia, China and the Middle East respectively. A hub market is not a data center park or campus (although these may form part of the hub) since these are local development sites, not complete economic entities in their own right.

Changes in the global configuration and role of data centers have altered the profile of the hub.

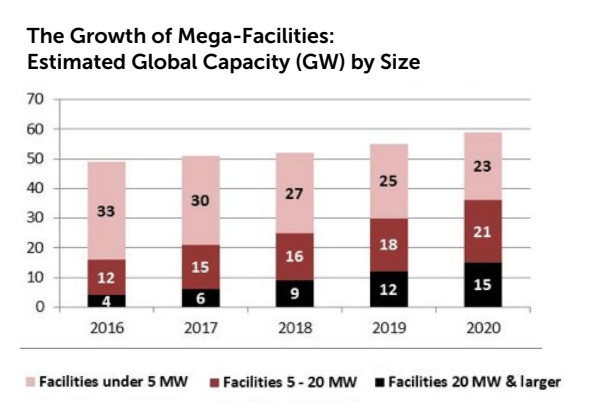

The global cloud providers with their hyperscale facilities are able to establish hubs well away from densely populated areas in order to benefit from available space, sustainable sources of energy, connectivity and incentives for investment. The growth of very large, power-hungry hyperscale data centers to support the growth in demand from colocation, cloud and other data center services has created significant growth in markets such as Ireland, the Nordic countries and US states in and adjacent to the Rockies. Thus the location of the hub for the hyperscale data center has moved from the urban to the outer city and then to the wilderness, and this trend will continue as these mega-facilities increase in number. This change in part has been enabled by the increased number and reach of fiber connections. Interconnectivity – access to multiple network providers, to the multi-cloud, to dark fiber, cross-connects and meet-me rooms etc. – is now enabling hubs well beyond urban areas.

Yet the role of the traditional hub market will not disappear. Hub markets are urban in order to provide the necessary economic momentum to support a data center sector, and it is no accident that hub cities tend also to be economic hubs at the intersection of trade routes. Local, urban data centers will be required to house highly latency sensitive applications, such as synchronous replication; and also where a particular location is required for compliance with data residency regulation.

However, hub cities tend now to be following the patterns established by other urban centers. The growth of data centers in London has been outside the M25 along the motorways connecting to the rest of the UK while the asset base in New York has shifted from the five boroughs to less densely urbanized areas of the Tri-State. There are various reasons for this – the availability and price of suitable space; sites that are easier to secure and protect outside the city; legislation that may restrict central city data center development (such as storage of diesel, for example); greater competition for key resources (power, water, IT skills); higher air pollution; older buildings and infrastructure, and, increasingly the fact that housing a full data center (as opposed to a micro or edge processing unit) in a central city location is unnecessary.

Hub cities have therefore developed hinterland zones and not necessarily in the same legislature. For Hong Kong, this is the Pearl River Delta in Mainland China while major developments in Benelux have moved outside the Amsterdam urban area. Singapore is supported by the Iskander development zone in Johor, Malaysia immediately to its north. It is likely that in the future hub cities will need to develop or have access to Real Estate in which to expand in order to maintain their premium IT position. It is obviously important for the hinterland to offer a similar quality of design and operation to the hub, to offer suitable connection into the hub and a legislative environment consistent with it in order that the two sites can offer suitable levels of integration and latency for the transfer of data, for back up, for service activation, for network security and for efficient portfolio management.

In the longer term, the role of the hub city and the remote hyperscale data centers will be reinforced as data processing moves towards the edge of the network. The 'edge' will tend towards the greater number of people and connected devices in urban areas while the core cloud facility will remain away from the data sources. The proximity of the core facility will not be an issue since one of the key drivers of edge computing is to reduce latency by processing and acting on data locally and immediately and as network speeds evolve with the development of technologies such as silicon phonetics so the requirement for urban hubs to be supported by connected hinterland zones will increase also.

We made some servers - think you can do better? Let us know @DCDnews

Re-shaping data center servers

They looked solid and unchanging, but now servers are being remolded by the creative minds of the industry, says *Dan Robinson*

Dan Robinson
Freelance

Servers are the linchpin of the modern data center, but the server market is in flux at the moment as trends such as cloud computing, social media and analytics are changing the demands on compute power, plus there is a growing need for greater energy efficiency and flexibility.

The past year has seen something of a slowdown in server shipments worldwide, with recent data from IDC indicating a decrease of 3.5 percent to 2.55 million units in the fourth quarter of 2016, while Gartner said that worldwide server shipments grew by just 0.1 percent during the entire year.

This has been attributed to reasons such as economic uncertainty and some customers re-evaluating their hardware provisioning criteria. Another factor is the uptake of public cloud services, with some organizations moving service provisioning to the cloud instead of running them on their own systems, leading to them requiring fewer servers while the large service providers operating those clouds need more.

"The dynamic of hyperscale being the segment of customers that are doing more purchasing and building more capacity is true to an extent," said Adrian O'Connell, research director for data center infrastructure and management at Gartner.

However, he added that enterprises currently still represent the largest part ▶

of the server market, and the move does not indicate a mass migration to public cloud, more a gradual drift.

"It's more a case of choosing which workloads might be going into the cloud and looking at how to introduce more efficiency in terms of your on-premise infrastructure. So it's a steady shift into the cloud rather than a sudden swing from one to the other," he said.

These hyperscale customers tend to be running somewhat different hardware to other customers. The stereotype image of a data center crammed with row upon row of racks stuffed with standardized "pizza box" rack-mounted systems is not far from the truth, but the differences are on the inside.

"The hyperscale buyers like AWS, Azure, Google and Facebook, they tend to buy these more commodity-oriented systems, that have the extraneous features that you typically get in an enterprise server stripped out of them," said O'Connell.

This means discarding any frills such as embedded lifecycle controllers to deliver a bare bones system that is standardized and keeps costs to a minimum. However, this does not mean these are second rate, as they are typically stuffed with large amounts of memory and the latest processors.

Open rack designs are like blade servers done better: shared power and cooling, with no lock-in

Enterprise customers, in contrast, tend to operate a more mixed environment, and thus have a more varied estate of servers comprising rack and blade systems, and even some tower chassis servers.

Blade servers can trace their roots back to early attempts to pack more compute power into a given space. Typically, this is accomplished by tightly packing multiple server modules or 'blades' into an enclosure that provides the power, cooling and connections to the rest of the infrastructure. Each enclosure then mounts into a data center rack.

However, while blade systems account for about 20 percent of overall server spending according to Gartner, there is little or no standardization in their design. This means that enclosures produced by one vendor cannot be fitted with server blades from another vendor, leading to a risk of lock-in for the customer.

Meanwhile, another format that has been growing in acceptance over the past couple of years is density-optimized servers, which combine some of the features of the blade and rack approaches. These typically provide multiple server nodes in a 2U or 4U rack-mounted chassis, often with the flexibility to mix and match different modules to meet customer requirements.

For example, Dell's PowerEdge FX offers customers a range of compute and storage modules, including two-socket and four-socket server nodes, plus one that comprises four separate nodes in a single module for high density configurations.

However, it is the rack-optimized format that still accounts for the largest share of the server market, typically around half of all systems sold.

"The 1U, 2U, 4U rack-optimized systems, they are still a pretty big chunk of the market, so a lot of enterprises, and a lot of the more traditional service providers, the tier 2 and 3 service providers, are still buying quite a lot of rack-optimized systems," O'Connell said.

This includes hyperconverged infrastructure (HCI) systems, which are a small but growing part of the overall server market. These are appliance-like systems that use internal direct-attached storage in a cluster of nodes to create a shared pool of storage, and are designed to scale by simply adding more nodes.

Hyperscale customers, however, are able to use their bulk purchasing power to specify custom designs that meet their exact requirements. The Open Compute Project (OCP) started by Facebook, for example, provides a set of specifications that large customers can take to an original design manufacturer (ODM) for manufacture.

A similar scheme, Open19, has been started by LinkedIn, aimed at meeting the needs of smaller operators and enterprises, while Microsoft has developed its own specifications that are also available via the OCP.

Ironically, some of the specifications produced by the OCP have been likened to blade servers but done better, because they stipulate shared power and cooling for all the systems in a single rack, while the specifications are not tied to one vendor.

However, the vendors offering these often do not provide the same kind of pre-sales and post-sales support that customers expect of a traditional server manufacturer, according to O'Connell. Again, this may not be an issue for those hyperscale operators.

Finally, one of the surprise winners in the current server market is Huawei, which saw a 64 percent increase in shipments for the fourth quarter of 2016 compared with the same period a year earlier, according to Gartner's figures. This saw it overtake Lenovo to become the third largest server vendor worldwide, thanks to aggressive pricing, but also its existing sales channel partnerships.

"Huawei has the advantage of its networking business and the routes to market it has established with that, and the reputation it has thanks to those, and proven capabilities around support and after-sales service," said O'Connell.

ARM and Power: taking on Intel-based servers

Data centers are currently dominated by x86 systems, but this architecture has been seeing renewed challengers of late, from chips based on IBM's Power and the ARM architecture.

The formation of the OpenPower Foundation opened up IBM's architecture to enable partner vendors such as Tyan, Supermicro and Wistron to build and sell their own systems based on Power processors. These are claimed to be able to deliver higher performance for demanding applications than Intel-based alternatives, as well as offering more performance at a given price point, according to IBM.

Meanwhile, Microsoft surprised everyone by showing off Windows-based ARM servers at the OCP Summit in March. These are for Microsoft's internal use only, but demonstrate that it is feasible to run Microsoft's cloud services on ARM systems.

Gartner pours cold water on the notion of any upset happening soon, however.

"When it comes to ARM or OpenPower, our position for a good few years has been one of there being a lot of potential, but it being no more than potential right now," said O'Connell.

"Is the price, performance, or power issue sufficiently compelling? It isn't good enough to just have parity with [x86], there has to be a significant reason for users to move their infrastructure away to this alternative thing, and we're not convinced there is this compelling benefit yet," he added.

Dreaming of Net Zero energy

Peter Judge
Global Editor

DCD Zettastructure, London, 2016

Can data centers cut their reliance on the grid, or even become net contributors? *Peter Judge* heard the arguments

Data centers have been cutting their energy for some time. Can they go off-grid and become independent - or even net contributors to the grid? Last November, DCD Zettastructure hosted a debate - and here are the highlights.

The idea isn't new: "A long time ago I surveyed 100 data centers for Enron, with a view to making them contribute power back to the grid using their generators," said energy expert Professor Ian Bitterlin. "It turned out not to be practical then, because the generators weren't rated for that job, and the company didn't want to use its generators to help the grid - when the generators were there to protect them *from* the grid."

That point of view is still widespread, but now the idea has been proven. The generators at Microsoft's Cheyenne site are run by the local utility and serve both Microsoft and the surrounding community.

The need is increasing. "We're switching off coal, which is 30 percent of our supply," said Emma Fryer, of trade body TechUK.

Utilities in various countries offer demand reduction schemes, whereby a facility can elect to switch over to its backup system at times when the grid is likely to be heavily loaded. The benefit is that extra capacity is not needed, and wasteful fossil base load generation is not required.

"The days of the passive consumer are coming to an end," said Russell Park, of the Distributed Energy team at Centrica (the parent of British Gas). The Centrica team is "dedicated to developing demand side response and distributed energy," he said. "We've diverted some £780 million ($973m) from generation into distributed energy."

Data centers are key to this, because they have high-grade generators ready to use, which could significantly help the grid. However the facilities are also critical, so they will not do anything which poses significant risk to their services.

"We spend between $30 million and $40m on electricity in EMEA every year," said Doug Balchin, director of critical infrastructure and utilities at network firm Level 3." At the moment we have about 60MW of backup power systems."

To take part in these programs may require an investment in more suitable generators, said Russell: "This is also an opportunity to invest in those assets and pay for that investment through a return on either revenues or cost savings, if a generator is coming towards the end of its natural life." The scheme can have a payback of three years or less, he said.

Running generators with a load occasionally is also good for them, so the schemes can generate other benefits for participants. "Running them on load is good for them, running off load is bad for them," said Bitterlin.

However, there are risks: "If you're running on the grid you're running through a transformer back to the grid," Bitterlin said. "If you have a short circuit event you may lose your data center."

Russell, agrees: "There are very, very few investments without risk, and I'm not suggesting for a moment that we risk the operational integrity of anybody's data center. The due diligence that we would undertake would ensure that the work we're doing would not impact on a backup generator. That's priority one." He says the data center operator keep control over the generators.

> *"A search engine might take the risk of demand reduction; for a bank the reward is negligible," Prof. Ian Bitterlin*

"The customer gets paid just for having it available," he said. "If we do use it we can operate it from the virtual power plant which we have in our headquarters in Windsor."

Using the on-site generator can allow companies to take part in the utility's Triad scheme, which affects pricing based on the demand in the three busiest half hours of the year, said Balchin. "Our colleagues are already operating the scheme in North America, so it's not a brand-new concept for Level 3. We're expecting to run our engines no more than 70 hours a year."

The actual value of such schemes needs close scrutiny, said Prof Bitterlin: "The data center never runs at full load so unloading the data center from the grid to reduce the grid's load only frees up a proportion of the facility's capacity. You may have 10MW of generation on site, but if your load is only 4MW then taking yourself off the grid is only giving the grid 4MW of relief, whereas running it in parallel with the grid and back-feeding onto the grid gives the grid 10MW."

A lot of this has been a phony war so far. Some data centers in London have belonged to similar schemes for years, and never been called on to donate their power. "They've had rewards and never had to take the risk because the global rumors of the grid's death have been exaggerated for quite some time," said Prof Bitterlin.

In the end, each possible participant has to assess the risks and rewards, said Bitterlin: "I can imagine a search engine thinking the risk-reward was well worth paying for because they didn't even care about the 20-minute outage, whereas a bank would think the reward was negligible in relation to the risk."

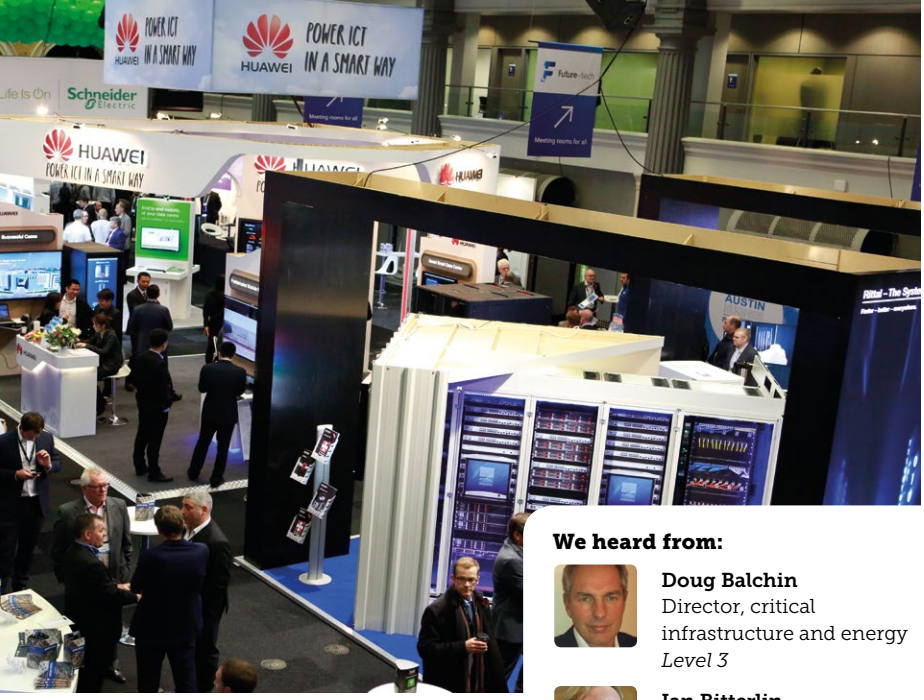

We heard from:

Doug Balchin
Director, critical infrastructure and energy
Level 3

Ian Bitterlin
Consulting engineer & visiting professor
Critical Facilities Consulting

Andrew Donoghue
European research manager
451 Research

Emma Fryer
Associate Director
techUK

Russell Park
Head of energy & sustainability solutions
British Gas

Regulations may have to catch up with the scheme, however, said Fryer: "The minute you move to elective generation, you start following foul of air quality control legislation." Running diesels will be seen to be a bad thing, even if they are feeding the grid and avoiding other emissions elsewhere but using them for disaster recovery is allowed, and feeding the grid is seen to be a matter of choice: "I haven't won that argument with [the regulators] yet."

The argument isn't a done deal either, because small generators on data center sites are arguably a lot less efficient than the plant in the centralized stations they displace, said Bitterlin: "When a diesel engine burns fuel oil, it's 33 percent thermally efficient. When you burn gas in a combined cycle turbine it's nearly 60 percent efficient."

We can still hope for a technology solution, according to analyst Andrew Donoghue of 451 Research: "Battery storage is now becoming far more viable. We've done a lot of evaluations and it's now a very compelling argument to attach additional batteries to balance the grid."

Although batteries are still expensive compared with the amount they can store, they could level some peaks of demand, and even a small difference there could reduce the generation capacity required.

"Storage is the Holy Grail," said Balchin. "Solar power is fantastic during the nice sunny day, but not in November cold evenings."

Another end-game might be small edge facilities, said Bitterlin: "It's easy to foresee a time when data centers don't take any energy at all. It's simply 'follow the edge.' Use 100kW data centers, never any bigger, and embed them in hotels, hospitals, and office complexes around the city. Take all the waste heat direct from the chips and use it. A hotel can absorb 100-150kW of hot water on a 24 hour basis. The energy goes in, the hotel gets free heat, and the mini edge computer center runs at zero cost."

It has to be edge, because heat can't be transported, he said: "The grade of heat is so poor that you can't export it further than a hundred yards."

The answer will involve a combination of all these ideas. "There's a palate of technologies which are available to all of us, and every situation requires a different hue," said Russell. "I think there is an uncomfortable marriage between government, the National Grid and the suppliers - if you can have a marriage for three people. But we know we have got to come up with a solution."

The Big Discussion on data center energy took place at DCD Zettastructure in London, Nov 2016.

DCD > Webscale | San Francisco

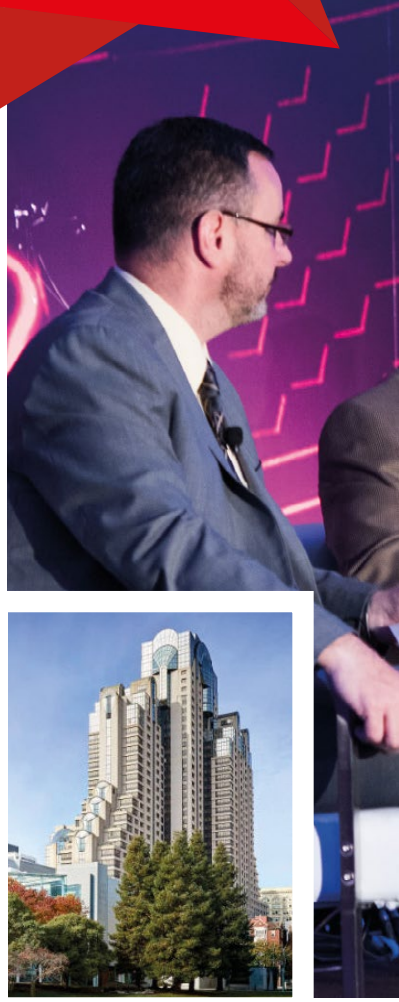

EVENT PREVIEW

The digital factory of the future will be Webscale and EnergySmart

June 19-20 2017 // **Marriott Marquis**, San Francisco

In the zettabyte-scale era, the data center is the new engine of the economy and electric power becomes the new oil

Marriott Marquis Hotel, San Francisco

Hyperscale digital infrastructure has got a new name - business scale - to best describe the ubiquity with which webscale technology architectures touch nearly every aspect of modern economic life. Digital infrastructure transformation comes with a transformation of electrical energy - increasingly clean, carbon-neutral, smart, and automated. These two realms of infrastructure systems are profoundly intertwined.

The two-day DCD>Webscale Conference and Expo, June 19-20, at the San Francisco Marriott Marquis has two central ideas: webscale cloud enterprises like Google, Microsoft, Amazon, HPE, IBM, eBay, Facebook, Salesforce, LinkedIn, Microsoft, Oracle and others will drive the new economy; and smart, digital "Energy 4.0" will make it possible.

"Every business today is somewhere on the webscale transformation journey," says conference chair Bruce Taylor. "Some are already truly digital-first, Internet-facing, network-edge and cloud-native. They are increasingly software-defined and data driven; incorporating AI and machine learning into both their external business development and their internal, full-stack infrastructure management.

"This conference is all about where this class of enterprises is headed, and about what others at an earlier stage in the digital transformation of their businesses can learn from the leaders."

+ Transforming the world of making stuff

Day 2 - the Webscale Summit on Tuesday, June 20 kicks off with a bang. Autodesk is iconic in architecture, engineering and construction, from data centers to giant airplanes and beyond. or just about anything else, than is Autodesk, and today the software giant that serves just about everyone who makes stuff, is bringing AI and machine learning in the cloud to its global customers.

There likely isn't a layer of the entire digital infrastructure cake that Autodesk's advanced tools and generative design data doesn't touch, transforming the design/build experience with AR/VR immersive visualization of generative 3D building design/build computation.

Lloyd Taylor is Autodesk's senior executive responsible for the company's cloud infrastructure, that delivers the "Future of Making Things," moving data from design through to implementation. "It's the application of industrial IoT to products or building operations or through to the actual manufactured product, such as the iPhone," says the computer scientist. "We have the ability to run full product and building ops simulations through AI/ML techniques beginning at the design phase."

Following Lloyd's keynote, data center journalist Rich Miller, founding editor of *Data Center Frontier*, leads a panel of AI and machine learning experts on how this high-form of software-defined smart technology is being applied to transform the modern hyperscale cloud data center beast.

IBM, Splunk, CBRE, Schneider Electric, Intel, LinkedIn, Microsoft, Spotinst, Rittal, Google, NetApp Solid Fire, TicketMaster/LiveNation, Switch - These are just some of the brand names across the full spectrum of the digital infrastructure ecosystem bringing their A-game expertise to the DCD>Webscale conference and expo to share their knowledge of how the hyperscale cloud factory and will evolve toward the future of the global economy.

University of Southern California's Dr. Julie Albright, who stunned the opening keynote audience for DCD>Enterprise in New York in March with her revelations about the impact of social and mobile media on society and culture, will lead a number of conversations over the course of both days.

"That's one of the most profound talks I've heard in a very long time," said Christian Belady, who leads Microsoft's global infrastructure and strategy, after Albright's New York keynote.

+ Energy 4.0 - Clean, smart, secure, always-on

Day One, Monday, June 19 is dedicated to the digital infrastructure ecosystem of electrical energy - the "oil of Industry 4.0" - without which the digitization of the global economy is impossible.

The electrical energy industry is undergoing the same scale and complexity of smart, digital transformation as is the data center itself.

DCD>EnergySmart is convening a full-afternoon collaborative leadership conversation, bringing together the primary energy and power stakeholders for the new economy. "For the first time ever, we're asking the people who carry the weight of responsibility for the future," says George Rockett, DCD's founding CEO. "We want to gather a broad representation of the executive and technical professional leadership from the supply side and the demand side of the new industrial economy, along with policymakers and regulators to explore common smart technology interests."

What do data center power engineering and management professionals need to know and what can they offer to the electrical energy community? The roundtable topics will cover power utility generation, smart energy network distribution, no-carbon renewables, regional and national smart grid transmission, microgrids, co-generation, demand response, storage and battery advancement, through the facility meter to the power path within the data center.

"For the purposes of this conversation, we're dividing the electrical energy universe into two segments," says Taylor. "North of the meter is everything on the public and private electric utility generation/transmission end; and south of the meter includes everything in the downstream power path within the data center." It's in this full path, he says, that smart power data technologies, such as the industrial IoT, big-data and predictive analytics, are now being applied to improve power reliability, security and efficiency. "Cloud services giants like Google and Microsoft are like big petri dishes for what's possible in wringing carbon out of the power grid in the time of climate change," he says.

Helping to bring together the knowledge resources for this DCD>EnergySmart leadership roundtable series are such organizations as the Electric Power Research Institute (EPRI), The Green Grid and the Business for Social Responsibility (BSR) sustainability organization and others.

EnergySmart is a limited-seating invitational executive roundtable series. Apply now to join it.

+ IoT Talent Consortium and the Infrastructure Masons

Two other influential groups will co-produce workshops on June 19. In the afternoon, the IoT Talent Consortium is an industry group including Cisco, General Electric, Global Knowledge, Microsoft, Rockwell Automation, Disney, Microsoft, MIT Sloan School of Management, Pearson Workforce Readiness, New York Academy of Sciences, the State of Illinois, and IQNavigator. IoT Talent's sole focus is on discovering and developing the talent that the IoT industry will require in the future.

Faculty for this unique IoT career workshop includes MIT Sloan School Associate Dean Peter Hirst; Rockwell Automation's Connected Enterprise development director Elizabeth Parkinson; Gordon Feller, director, Cisco HQ; Patty Burke, director of innovation and leadership solutions, Center for Creative Leadership.

Registration for IoT Talent also gains all tutorial participants access to all of the DCD>Webscale conference on June 20.

Infrastructure Masons, founded by Uber Compute's Dean Nelson, will meet with Silicon Valley data center executives for brunch to consider big future challenges to the digital infrastructure industry.

www.**dcd**.events

> Community

"The dilemma for all IT organizations is not whether to change the way they do things, but which of the new paths to take. [This] conference program embraces the need for a much more holistic discussion about the future of data center and cloud infrastructure that I am sure will help people make better informed decisions."

Kfir Godrich | BlackRock

Events

> **Middle East** | Dubai
May 9 2017

Upgrading the region's critical data infrastructure for IoT, cloud and smart cities

Training

Data Center Cooling Professional: London, UK
April 19-21 2017

Data Center Power Professional: London, UK
April 24-26 2017

Data Center Design Awareness: London, UK
April 24-26 2017

Training

Data Center Design Awareness: Lima, Spain
April 17-19 2017, TBC

Data Center Cooling Professional: Madrid, Spain
April 24-26 2017, TBC

Events

> **Webscale** | San Francisco
June 19-20 2017

The global summit for webscale infrastructure builders

Training

Data Center Design Awareness: New York, USA
April 24-26 2017, TBC

Events

> **Argentina** | Buenos Aires
April 25 2017

The congress on digital infrastructure and cloud

> **Colombia** | Bogotá
June 14 2017

Training

Data Center Design Awareness: Brasilia, Brazil
April 10-12 2017, TBC

Every year we can see that more senior people are attending this event. The way the program is structured and the fact that is covers so many different topics at the same time makes the event even greater.

Sinem Cantürk | KPMG

Simply the best data center conference in the Middle East.

Dick van Bladel | IBM

Events

> **Indonesia** | Jakarta
April 6 2017
IT transformation and hybrid architecture

> **Focus On** | Hyderabad
April 27 2017
The future of India's digital infrastructure

> **Malaysia** | Johor Bahru
May 18 2017
Data center and cloud for the new digital economy

Training

Data Center Design Awareness: Jakarta, Indonesia
April 17-19 2017, TBC

Energy Efficiency Best Practice: Jakarta, Indonesia
April 20-21 2017, TBC

Data Center Design Awareness: Bangkok, Thailand
April 24-25 2017, TBC

Critical Operations Professional: Bangkok, Thailand
April 26-28 2017, TBC

It was a great event in 2016 in a venue that creates much more intimacy and even better customer interactions. The content was fantastic with some great speakers from blue chip organizations. All in all, a great learning and thought leadership event.

Matthew Baynes | Schneider Electric

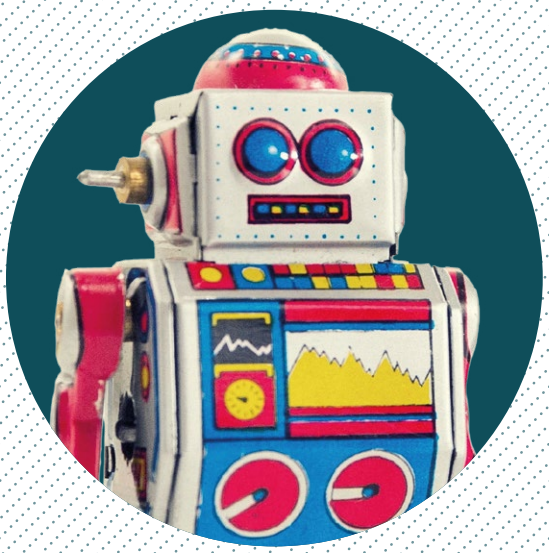

Get these humans out of my data center!

Humans were never meant to work in a data center environment. They are soft, fragile and susceptible to electric shock. They forget best practices, make mistakes and invite security risks. They are driven by emotions like anger, greed and a desire for revenge. And yet we still design data centers around people, not around technology.

If we could eliminate humans from our facilities, infrastructure would become more reliable and more efficient. We could raise operating temperatures from the recommended range of around 20°C to 25°C (68-77°F) all the way up to 45°C (113°F), resulting in considerable savings on cooling costs. Several vendors have previously confirmed that their hardware can take this kind of harsh treatment. We wouldn't have to care about air quality or humidity, or even keep the lights on.

Most importantly, we would eliminate human error – one of the main reasons for unplanned downtime. According to research by the Ponemon Institute, 22 percent of all data center outages in 2016 were caused by people, not hardware.

> *"Look at you, hacker: a pathetic creature of meat and bone, panting and sweating as you run through my corridors. How can you challenge a perfect, immortal machine?"*
>
> *SHODAN, System Shock*

There are plenty of organizations that are taking important steps towards automating their infrastructure. For example, Microsoft Research continues its work on Project Natick, sinking autonomous data centers to the bottom of the ocean. AOL claims to have operated a 'lights out' data center called ATC since 2011, but details about the facility are few and far between.

Both IBM and EMC previously equipped iRobot's Roomba – the autonomous vacuum cleaner - with temperature sensors in order to build heat maps of their server farms, representing one of the early uses of actual robots in a data center.

Sony's Everspan Library System is another example of data center robotics: this modular appliance holds thousands of high capacity Archival Discs that are managed by an automatic arm, like a futuristic jukebox. Meanwhile American startup called Wave2Wave has developed Rome, a series of switches that enable physical fiber connections to be made automatically.

Software automation driven by tools like Chef and Puppet and trends like SDN are also doing their part to help minimize the need to interact with actual hardware.

And there's no reason why we couldn't completely eliminate people from the equation: after all, a data center is nothing but a warehouse full of servers, and warehouse automation is something that is being actively investigated by several retail giants, with Amazon emerging as one of the leaders in this field. We will always need people to look after the infrastructure, but soon, a time will come when we no longer need to share the same space.

Max Smolaks
News Editor

DCD

> Software-Defined | Supplement

Powered by

JUNIPER NETWORKS

INSIDE

A Special Supplement to DCD April/May 2017

Powered by

Contents

Features

Opinion

A question of definition

Has the software-defined data center (SDDC) arrived yet? I'm afraid that depends on what you mean by an SDDC. Like so many other concepts in this industry, the idea of a software-defined data center can mean several different things - and some of those things aren't really here as of yet.

Deep down, it's all about automation. IT services have been made simpler in the cloud, but SDDC promises to automate their delivery, so they can be provisioned and maintained without human intervention.

To this end, we need software-defined storage, software-defined networking services, and software-defined compute. All these should be available as pools of resources, to be deployed with a few clicks.

We're making progress towards that (p4). Platforms that offer these pools are emerging from companies such as HPE, as well as from the collaborative efforts going into the OpenStack project.

The open source options may be less mature and more complex, but the proprietary routes may, as always, involve possible lock-in.

Further up the stack, we come to the levels where users encounter the SDDC (p10). Here, we won't argue about the merits of various packages and protocols.

The question will be whether a given SDDC approach is practical. Does it allow a service provider to meet customer needs better, perhaps by offering new services?

This, far more than the technical aspects of the solution, will be the arena where SDDC makes its mark, although we suspect users may be unaware of the impact of the technology: the concept and the abbreviation of SDDC is far too nerdy to ever be a consumer brand - even for the kind of consumers we meet in this market.

Networks have been the test case in many ways. Virtualization had taken hold in the server space, and storage was acquiring levels of abstraction that effectively turned it into a pool of capacity.

Networks seemed to be stuck with proprietary systems, and uncessary hardware, till software-defined networking (SDN) separated the functions from the the underlying hardware, created the "software-defined" moniker, and now lead the way (p12).

Physical infrastructure can be overlooked in all this, but a data center is not software-defined until all its resources can be managed and controlled by software.

DCD's Bruce Taylor (p14) says that the software-defined data center is little more than a sham if it only manages the IT stack, north of the rack. It needs to extend its power south into the mechanical and electrical infrastructure if it is to deliver its promises.

Get that part right, and the whole facility is defined by software and delivered as an easily consumed service.

As Bruce suggests, this is when software finally eats the data center.

Peter Judge
DCD Global Editor

Dan Robinson
Freelance

Automatic for the data center operators

Everything is going software-defined, says *Dan Robinson*. And the main thing this means is automation

Over the past several years, "software-defined" has become another of those terms that gets used so much that its meaning becomes blurred. However, software-defined infrastructure is a key part of the way data centers are currently being reshaped in order to meet the changing requirements of modern IT and applications, which are now more distributed and dynamic in nature.

One of the stumbling blocks that enterprises and data center operators face is that IT infrastructure is increasingly complex to deploy and configure. Applications often depend upon an array of other services and resources to be in place before they can function properly.

Solving this problem calls for a certain degree of automation, according to Andy Buss, consulting manager for data center infrastructure and client devices at IDC Europe.

"Any business that is on a journey to digital transformation needs automation, as automation is the ability to put actions into IT policy," he said, adding that "moving from dev to ops is about being able to deploy [infrastructure] as easily as possible."

The implication of this is that the infrastructure must be capable of being reconfigured or repurposed using software, rather than having engineers go into the data center itself and physically reconfigure the infrastructure.

This kind of approach is already used extensively by many of the biggest Internet companies, such as Google and Amazon ▶

Getting the Synergy right

HPE's Synergy infrastructure platform, which the firm began officially shipping in January, is one example of what infrastructure designed from the ground up to be software defined looks like, although HPE does not refer to it as such, preferring to call it "composable infrastructure."

The big idea behind Synergy is that is "defined by code," according to the firm, meaning that it is designed to be almost entirely driven by templates that specify the resources required for specific applications and services, and how these should be configured.

"It's a stateless computer, designed to be a reconfigurable pool of resources," according to Buss. It can be regarded as a set of storage, networking, compute and memory resources that can be pulled together to create a virtual computer for whatever task is at hand.

The key part of Synergy is that HPE's updated OneView management system can provision and manage the bare metal, so it can be used to stand up a database cluster running on dedicated hardware as easily as provision virtual servers and applications running in containers.

The caveat is that Synergy is effectively a combination of special hardware and software that is available only from HPE. While customers can choose from a range of compute and storage nodes to suit their requirements, they cannot mix and match with third party kit within a Synergy deployment.

Web Services (AWS), as it is a fundamental requirement of operating a public cloud platform, where services may be continuously starting, scaling up, and eventually releasing resources again when they terminate.

In other words "software-defined" refers to infrastructure and resources that are decoupled from the hardware they are running on, and are able to be dynamically configured and reconfigured under software control, typically using application programming interfaces (APIs).

Among the first components of IT infrastructure to get the software-defined treatment were servers, through virtualization. A virtual machine running in a cloud is effectively a software-defined server, since it has been decoupled from the underlying physical hardware, and can be moved from one physical server to another in the event of a hardware failure, or for other reasons such as load balancing.

The past few years has seen the rise of containers, which are less demanding of resources than virtual machines. Platforms such as Docker also have more of a focus on distributing and managing applications as a collection of containerized functions.

As trends like cloud computing have become an increasingly important part of the data center, the move to make more and more of the infrastructure software-defined and therefore more agile has grown.

However "software-defined" can mean different things when the term is used by different vendors, and there is not always a widely agreed standard definition.

For example, software-defined storage can include storage virtualization, such as EMC's ViPR platform, which enables a customer to build their storage infrastructure using a mix of arrays, including those from multiple vendors, and manage it all from a single console.

But software-defined storage more commonly refers to software such as the open source GlusterFS that runs on a cluster of commodity server nodes and uses these to create a scalable pool of storage.

This model is seen in hyperconverged infrastructure (HCI) systems, in high performance computing (HPC) clusters, and in very large Internet companies like Google and Facebook because it is easy to provision and scale.

$25.6bn
revenue from SDDC in 2016 (Markets andMarkets)

However, software-defined storage is only as good as the hardware it is running on, and may not be as reliable as purpose-built enterprise storage arrays. The assumption is that redundancy and other capabilities like data replication will be taken care of elsewhere in the software stack.

Software-defined networking similarly covers a number of different approaches designed to make the network infrastructure more dynamic and easy to configure as required.

One approach is to separate the control plane, or management part of the network, from the forwarding plane, typically the switch hardware that actually routes packets around the network. The idea here is to centralize control over network traffic, using tools such as OpenFlow, a protocol designed to provide a standard mechanism for this, which is supported by many switch vendors, such as Dell, HPE, and Cisco.

A different approach is to virtualize the network, allowing the creation of logical networks that use the physical network to move data around, but which do not necessarily have the same IP address scheme, and may each differ in their quality of service and security policies.

This latter approach is typified by VMware's NSX and the Neutron module in the OpenStack platform, and is vital for supporting multi tenancy in data centers, especially those hosting public cloud services.

Virtualizing the network not only enables new network connections to be created without having to physically move patch cables around, but in the case of VMware's NSX, it also enables greater oversight of network traffic, since the switching and routing capability is integrated into the hypervisor and distributed throughout the infrastructure.

These strands – software-defined compute, storage and networking – do

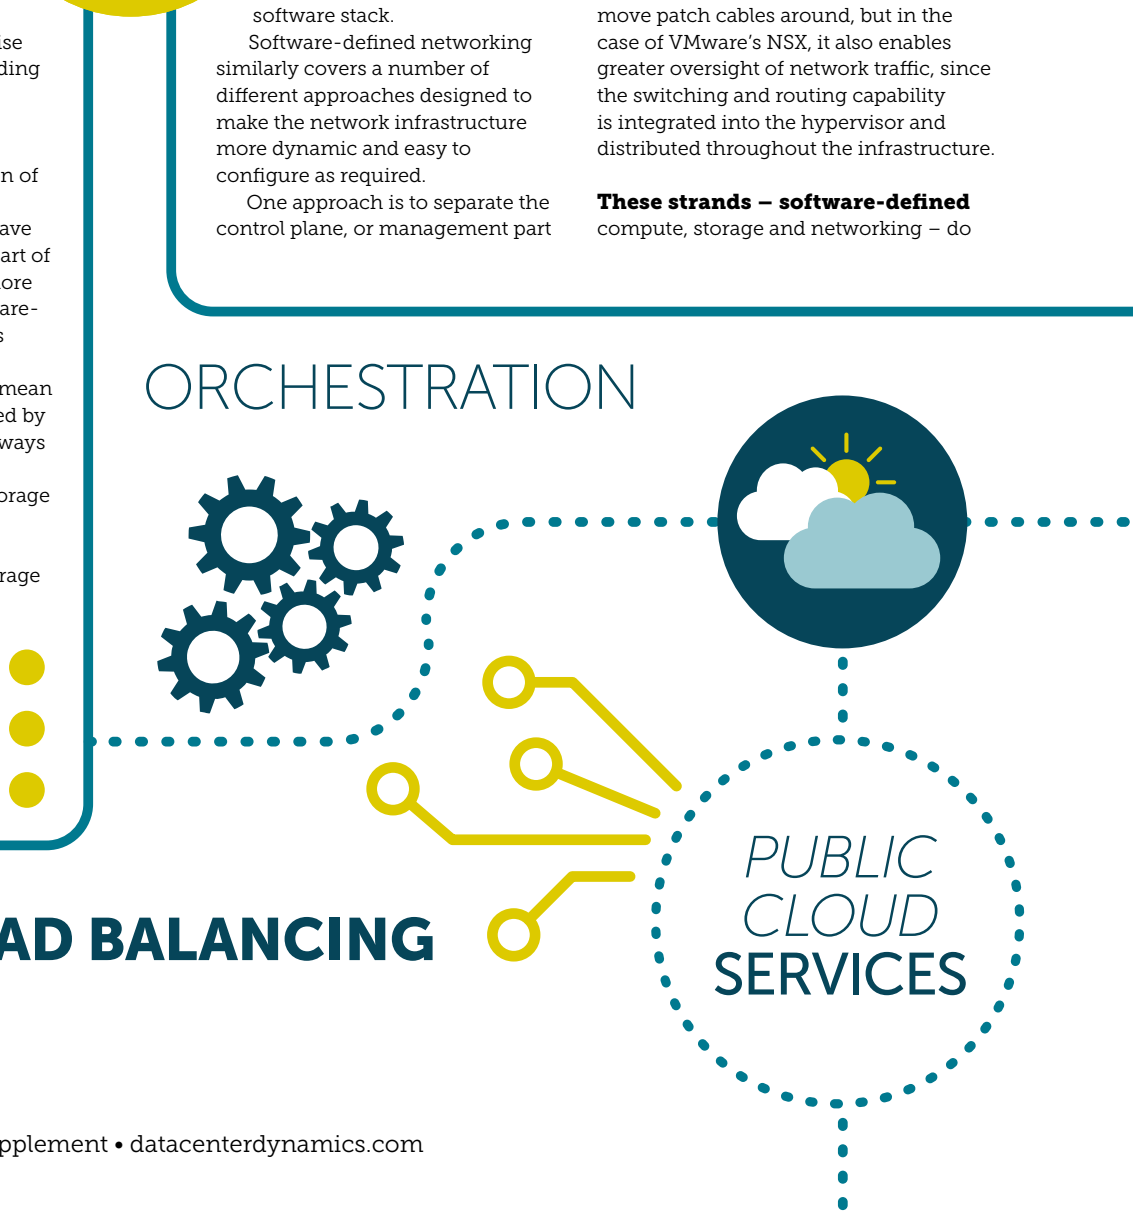

not exist separately, but are very often interdependent. Pulling them all together is the ultimate aim, in order to deliver what is termed by Intel and others as the software-defined data center (SDDC)

With SDDC, the entire data center infrastructure should be configurable under software control, making it easier to automate so that IT staff can devote less time to simply keeping the lights on, and this means that the fourth piece of the puzzle is management and orchestration.

"It's pointless having software-defined anything without automation and orchestration," said Buss. He cited tools such as Microsoft's System Center Operations Manager (SCOM) and VMware's vCloud Director as examples, but added that these have yet to reach the same level of maturity as those used by the big cloud providers.

Also less mature but gaining support among service providers is OpenStack, which is open and non-proprietary and thus independent of any single vendor, and which presents a set of APIs that can be used to plug in other software and services.

Overall, the SDDC approach has passed an inflection point and is growing in importance, according to Buss.

"The public cloud demonstrates it is feasible, and as we move to a multi-cloud world, there will be a need for compatibility between clouds which is driving a lot of thought in the industry now."

Open source SDDC

OpenStack is well known as an open source software project, widely used for cloud deployments. However, its strength lies in the fact that it is actually a management framework that is designed to be as open and extendable as possible, making it a general-purpose "integration engine" for operating IT infrastructure.

This can be seen from the number of telecoms operators that are turning to OpenStack as the linchpin for a modernization of their infrastructure, allowing them to replace costly specialized network hardware with software running on commodity servers that does the same job, an approach known as network functions virtualization (NFV).

One of the reasons for this is OpenStack's flexibility; it is now organized around core services including compute, storage and networking modules, while others are optional and can be used to provide services as required by the user, including telemetry, orchestration and even a database service.

Nova, OpenStack's compute module, is not tied to a specific hypervisor. It uses the Linux KVM by default, but also supports those from VMware and Microsoft, as well as container technologies such as LXC. Likewise, the Neutron networking and Cinder storage modules plug into a variety of other products and platforms, as in VMware Integrated OpenStack, where they are used atop VMware's NSX and VSAN, respectively.

If OpenStack has a weakness, it is that some aspects, such as automation and orchestration, are not as mature as in some proprietary cloud and data center platforms. However, thanks to its open APIs, developers and third party tools can fill in the gaps.

These APIs mean that users can reconfigure their infrastructure under programmatic control, a fundamental aspect of a SDDC.

OpenStack has a wide number of backers in the technology industry, including Intel, which has been using the software itself internally since 2011 to manage its own infrastructure, as well as contributing code to the project.

Containers open the cloud

App containers take virtualization even further, says Max Smolaks. Think of it as yet another level of software-defined abstraction

While virtual machines have done wonders for hardware utilization, a new approach to workloads is emerging that could enable data centers to use their resources even more efficiently.

Kubernetes and Docker are the most popular examples of systems that automate the deployment of applications inside software containers. They package apps and all of their dependencies into individual environments that all rely on the same Linux kernel, which means they require less overhead than VMs.

App containers are completely independent of the underlying infrastructure and can be easily moved across servers, data centers and cloud platforms. At the same time, seamless transition between development, test and production environments makes this technology indispensable for DevOps. And finally, due to their policy-driven nature and isolation from the host, applications packaged in containers are fundamentally more secure.

Containerized apps represent an important step towards microservice architecture – where various services are packaged separately and chained together using orchestration tools.

This approach is especially useful for massive software projects running in cloud environments, since it allows individual services within applications to be scaled up or down depending on demand. It also enables developer teams to work on software without having to be aware of what the other teams are doing.

Kubernetes was originally developed at Google while Docker was created by Solomon Hykes at the eponymous Docker Inc (formerly dotCloud). Today, both are open source projects with a massive following, which means they will continue to evolve rapidly in response to the demands of the early adopters.

Containers look like the future.

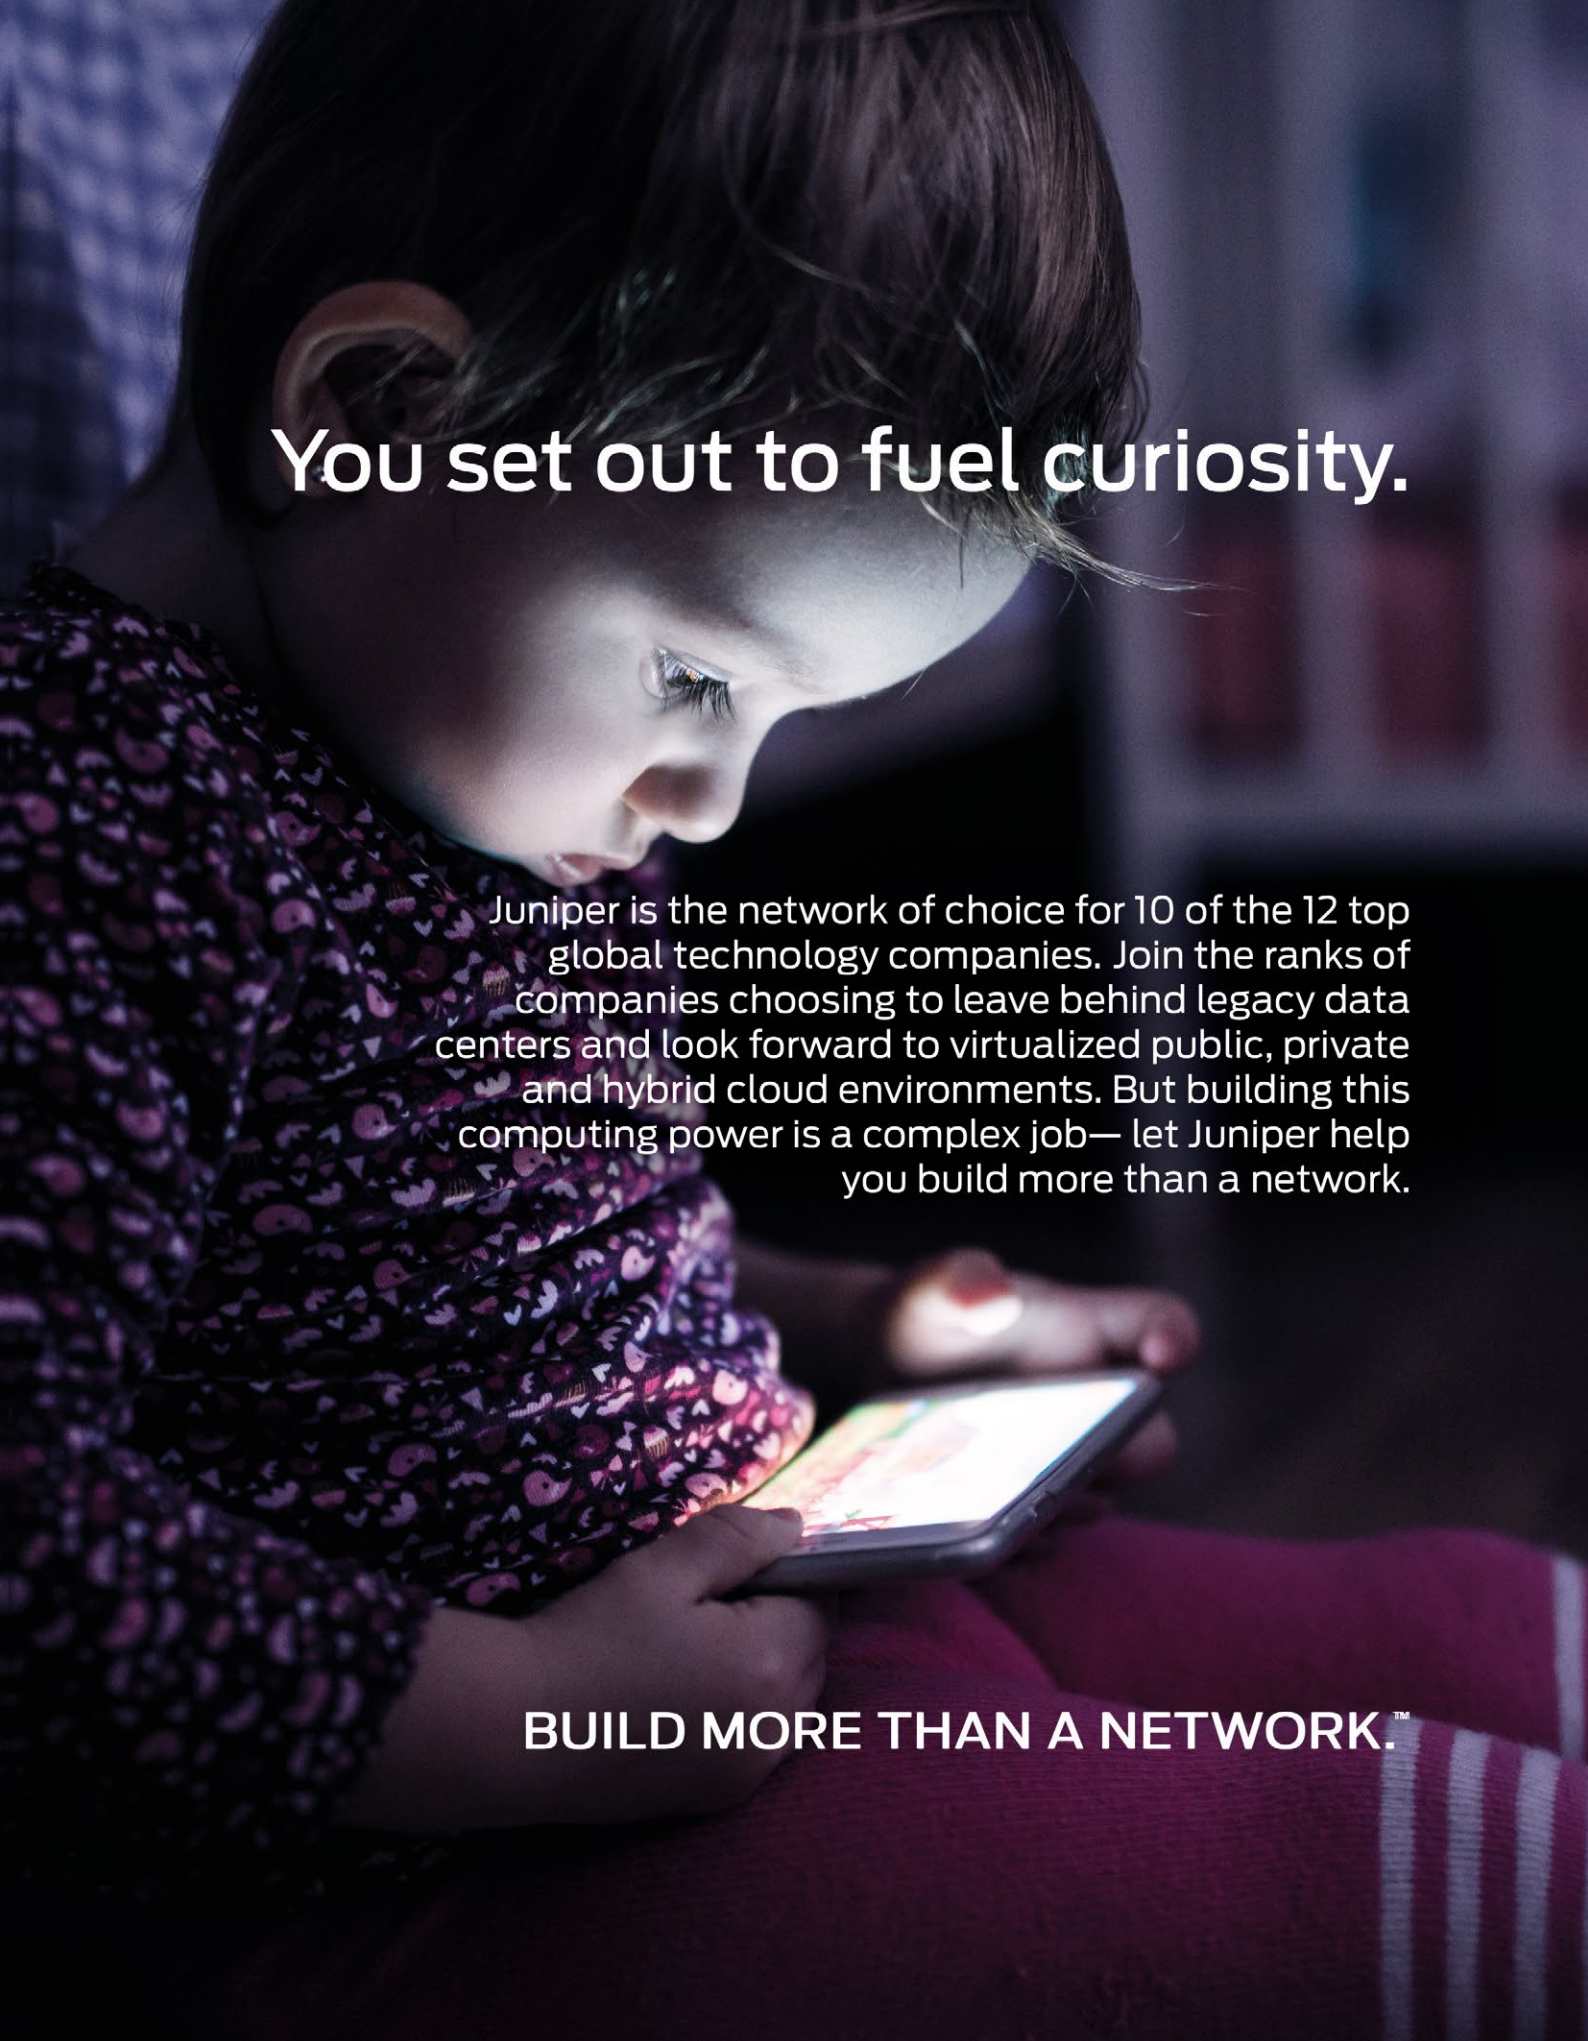
You set out to fuel curiosity.
Juniper is the network of choice for 10 of the 12 top global technology companies. Join the ranks of companies choosing to leave behind legacy data centers and look forward to virtualized public, private and hybrid cloud environments. But building this computing power is a complex job— let Juniper help you build more than a network.
BUILD MORE THAN A NETWORK.™

Software-defined? Let's aim for business defined!

Chris MacKinnon
North America correspondent

The real benefit of the software-defined data center comes when you connect it to a layer of business logic, reports *Chris MacKinnon*

Has the data center reached the pinnacle of automation? There is no doubt that facilities are more "robotic" and self-serving than ever before, but there could be more to come. Software is continuing its quest to become "the great definer" at the heart of the software-defined data center (SDDC), and the data center itself is taking on a more chameleon-like shape.

To Matt Altieri, marketing director at Device42, the SDDC means software-defined everything, where the idea of virtualization is now extended to all parts of the IT stack, resulting in delivery of infrastructure as a service (IaaS).

"Software-defined data centers are not only simpler and more manageable," Altieri told us, "but they also more easily align with companies' business needs. The speed of deploying systems is dramatically increased."

He says SDDCs have the flexibility to take on many configurations, capabilities, and forms, and can support companies ranging from the size and complexity of the FANG companies (Facebook, Amazon, Netflix, and Google) to companies with very much simpler operations.

The increase in simplicity realized from "converged infrastructure," as Altieri put it, paved the way for delivery of IaaS: Amazon Web Services changed the world of IT by pioneering a new way to deliver IT, and the other public cloud providers such as Microsoft, Google, and IBM soon followed. Altieri said these providers allow you to use application program interfaces (API) to create the infrastructure you need to run your business on demand. A SDDC differs from a private cloud, which only offers virtual-machine self-service, beneath which it could use traditional provisioning and management. Instead, SDDC concepts describe a data center that can span private, public or hybrid clouds.

Torsten Volk, managing research director with Enterprise Management Associates, and no stranger to the SDDC arena, believes that only in a SDDC can customers achieve policy-driven and fully-automated application provisioning and management.

Volk said: "Today, when the request comes in to host a specific enterprise application, IT teams (for storage, networking equipment and servers) have to crack open multiple hardware vendor specific command-line interfaces (CLI) and control panels to serve up the storage, network and compute resource pools needed by the virtual machine administrator to create the virtual application environment." But Volk said this is error prone, requires vendor specific skills, and can take days or weeks. He said these issues are the reason for the rapid growth of Microsoft Azure and Amazon Web Services, where line-of-business (LoB) developers already have fully-programmatic access to everything they need.

26.6%
predicted annual growth in SDDC till 2020 (MarketsandMarkets)

The "business logic layer" has always been Volk's vision of artificial intelligence/machine learning-driven (AI/ML) intelligence that can manage hybrid IT infrastructure in a manner that accounts for the business importance of each application, as well as for the organization's permanently changing strategic priorities. In short, the business logic layer turns the SDDC into the business-defined data center (BDDC).

Volk elaborated: "The crux is that we can only make the step to the business-defined data center if we make IT operations management sensitive to the business context, instead of rigidly enforcing best practices for application management. For example, in a business defined data center, the business logic layer makes the painful decisions in terms of which tradeoffs have to be made for IT to be optimally aligned with the business."

This could mean accepting the slow performance of a specific application as a price for optimizing the user experience of another, much more business-critical application. Volk said these decisions cannot be effectively made in today's siloed IT, as individual operators do not have access to the business context.

For Michael Letschin, field CTO at Nexenta, SDDCs are currently changing the entire data center landscape. In the past, facilities were made up using big block systems from legacy players, purpose-built to support a narrow, pre-defined set of workloads.

"This limited the end user's choice for technology and usage," Letschin said. "The newer software-defined storage model is shifting this paradigm and is giving organizations the ability to run an agile, scalable and cost effective SDDC."

Additionally, the selection of vendors for the data center has expanded and challenged the concept of "never getting fired for buying IBM," Letschin said. In today's SDDC, management and administrative staff are encouraged to look for innovative software-defined solutions on the market to address their data center pain points. He continued: "The benefits of transitioning to an SDDC also extend beyond the core technology to deliver significant space, power and cooling reductions. As a result of the clear business benefits that software-defined storage (SDS) provides, an increasing number of enterprises are shifting to the SDDC model.

According to Letschin, if the last five to ten years have been any indication, it will be almost impossible to predict the growth in the software-defined data center, but it will include integration with public cloud solutions to give even more "just in time" solution capabilities. He said: "We will see the integration of multi-data center capabilities. On the compute side, the rise of containers will bring the application back to the forefront and the idea of software-defined compute running on all full virtual machines will become obsolete. This increase in application based compute will lead to more automation, self-service and self-supporting infrastructure with AI taking a bigger role in actively managing the SDDC."

Letschin said storage will most likely undergo some of the biggest changes and, at the same time, remain the most consistent. Also, data integrity will continue to be key for compliance and security purposes. However, the use of commodity hardware will become more prevalent in the future with the ability to cram more capacity in less space while providing higher performance. "One thing that will start to become obsolete in the data center," Letschin said, "will be 'performance-only' solutions, and general-purpose storage will likely make a comeback because of the flexibility inherent to SDS."

Overall, Letschin says the data center will continue to operate as the nervous system of the business, but will be much more agile than the data center we recognize today.

What do the next ten years hold?

The concepts behind the software-defined data center have been dismissed by some as hype. Some critics believe that a minority of companies with completely homogeneous IT systems already in place (FANG companies, specifically) can truly transition to software-defined data centers.

The FANG companies have implemented this vision by investing years of engineering time and billions of dollars. But in our view, it is possible for any company to start on the road to ITaaS and reap their business-specific benefits. One company's approach to ITaaS will likely be different than that of other companies, because it will be built to support unique business needs. For example, Facebook's data center was created specifically to meet the demands of its product offerings; Cassandra sprang out of Facebook's needs.

Therefore, rather than look at what the FANG companies have done and try to duplicate their capabilities, start by understanding your own data center and your business requirements. The task is to figure out what you have and then make that more adaptable. You must understand how to best support the applications needed by your business so they can scale and adapt to your needs. If your data center is starting from a point of too much complexity, you'll need to simplify and pare down.

Matt Altieri, *Device42*

Peace, love and SDN

Software-defined networking is changing the role of switches, and the companies that make them, says *Max Smolaks*

Max Smolaks
News Editor

Virtualization has been a blessing for data centers – thanks to the humble hypervisor, we can create, move and rearrange computers on a whim, without thinking about the physical infrastructure.

The simplicity and efficiency of VMs has prompted network engineers to envision a programmable, flexible network based on open protocols and REST APIs that could be managed from a single interface, without worrying about each router and switch.

The idea came to be known as software defined networking (SDN), a term that originally emerged more than a decade ago. SDN also promised faster network deployments, lower costs and a high degree of automation. There was just one problem – the lack of software tools to make SDN a reality.

This was the hurdle faced by all networking equipment vendors, but those who saw that SDN would eventually become the norm, realized that success in this field would require a wide ecosystem of partners, even if they were also your competitors.

As with so many other areas of IT, the answer is in open source. If developments

> *"We're giving engineers tools to modernize how they manage networks"*

are shared amongst a wide community, then progress can be quicker as companies are not duplicating efforts within proprietary worlds. There is also a level playing field to compete on - but the best way to compete is to be on the teams building that playing field.

This was the scene in 2012, when network firms started buying SDN startups. Juniper Networks took the strategy seriously. It bought Contrail Systems - a secretive startup - on the cheap and quickly published its code under an open source license.

Today, the OpenContrail open source project team has drawn in employees from Nokia, Mirantis, Symantec, Canonical, IBM, AT&T and NTT Innovation Institute, and the supported version has created a new, thriving business within Juniper itself. One of the company's flagship products is Contrail Cloud - a mix of OpenStack, Open Contrail, Ceph and Puppet, bundled together with a host of minor enhancements.

When Juniper bought Contrail for $176 million, it had never shipped a single product. Less than a year later it released the code under the Apache license. This deal makes more sense once you realize that Contrail was founded by former Juniper employees who worked on the popular MX series routers and switches.

"People from Juniper left the company to go there, so there was always a relationship," Scott Sneddon, senior director for SDN and Cloud at Juniper, told *DCD*. "Even Kireeti Kompella, one of our early, really strong leaders in the development of MPLS [multiprotocol label switching] who had a lot to do with Juniper's success in its first fifteen years, he went to become a CTO at Contrail."

Those were the days when everyone was looking to buy into SDN. A famous example is VMware, which won a bidding war against Cisco for a higher-profile startup, Nicira, whose founders included Nick McKeown and Martin Casado, the researchers at Stanford who created the SDN concept. Nicira's SDN implementation is now in VMware's NSX.

Nicira was somewhat more expensive than Contrail: VMware ended up paying $1.26 billion. "Two months later, we sent a rover to Mars for just double that," Sneddon joked. By comparison, Juniper's acquisition of Contrail seems like a bargain.

Fast forward four years, and OpenContrail is an important part of both Contrail Cloud and Unite Cloud – the latter, launched in January, is Juniper's data center framework that aims to simplify creation and maintenance of hybrid and multi-cloud environments. It includes access to Contrail JumpStart service, which allows customers to experiment with open source tools.

$1.26bn
price VMware paid for Nicira

"Even if the customer doesn't adopt an SDN solution right away, we still have a way to have that conversation and help them evolve," Sneddon said. "I've been working on SDN solutions for about five years now, and I can honestly say that in more than half the meetings I go into, SDN is probably not what the customer needs. But even without an SDN solution on top of it, we have a really good automation framework: we've done a lot of work to develop a bunch of Python, SaltStack, Puppet and Chef interfaces for our routers and switches.

"What we're doing is we're giving the network engineers tools to modernize how they manage their networks. So I go and I talk to a customer about SDN, because that's my job: I lead with Contrail and we talk about the evolution and cloud platforms and systems. But if I'm talking to a hardcore network engineer who just builds data centers and doesn't have visibility into the applications that are running – that oftentimes goes over their head.

"So I'll start talking about how you operate it – because really, what we are trying to do with SDN is just automate the network provisioning process, and optimize how the network works for cloud platforms."

Contrail is meant to complement OpenStack, which means it's designed for cloud computing at scale. The centerpiece of the platform is the SDN controller, which defines how the network is laid out and what the topology is. It is accompanied by virtual routers that have to be installed on every server and linked to the controller. There's also an analytics component which monitors the state of the network, and a number of APIs for control functions.

"The physical network that's in place just becomes a transport layer that passes the packets over a routed network," Sneddon explained. "And we know how to manage routed networks very efficiently. And then we take the really complex service layer, the things that have to change when I start a new application, or spin up fifty VMs that span across a massive data center, and we build overlay tunnels to support those – and really, an overlay tunnel is just a VPN."

Open Contrail and Contrail Networking – Juniper's own version - are identical in feature set. Furthermore, Contrail Cloud releases are aligned with OpenStack community releases to maintain API compatibility and keep strong community support. The paid versions just add enterprise-level support and installation.

Even with the relative success of Contrail, Juniper is not content to rest on its laurels: last year, the company acquired AppFormix, a cloud management and optimization startup. It develops a server-level tool, suggesting that the software-defined approach could turn network companies into something beyond their old role as hardware merchants.

"You'll hear these cloud guys talk about the network, where the only awareness of the network they have is when it's broken. They don't care if you're using BGP or MPLS or VLANs, as long as it's there and it works. As a side effect, the purchasing decisions – which switch or which router they buy – sometimes become less important," Sneddon admitted.

"So for Juniper to be relevant in the future, we know that we have to play at a higher level. A lot of enterprise IT buying decisions are coming from cloud teams, not network teams. They have the dollars. So we feel like we have to develop a strong value proposition for cloud engineers and architects."

Software-defined has to get critical

— **Bruce Taylor** | EVP, DCD

You can't have a software-defined data center until it handles the mechanical and electrical parts of the facility, says *Bruce Taylor*

The true autonomous, lights-out data center, where the facility becomes a commoditized utility, may not be that far off, as the role of hardware shrivels and dumbs down. But has the ability to deliver the true software-defined data center (SDDC) been overstated?

A few years back, cloud was a pipe dream. Cloud required network, storage and compute to be combined in a cohesive, integrated, software-managed infrastructure. That demanded a highly abstracted (virtualized), automated, policy-based system that combined workload management, agile infrastructure provisioning, failover, disaster recovery and security. And that package simply didn't exist.

The virtualization and cloud pioneers knew they could pool IT resources, but in the early days they still hadn't really given a thought to creating physical data center infrastructure, including the guts of the critical environment: power and cooling (thermal management).

Now the SDDC is in labor after a very long pregnancy, promising to deliver a software- and data-led unified toolbox that presents the data center as an abstracted private cloud, available to multiple customers. But the critical environment is still not there.

SDDC is needed, because digital transformation is here. In 2016 the world will hit 1,000 exabytes (one zettabyte) of data traffic on the Internet, says Cisco's Visual Networking Index, with Internet data globally projected to grow at a CAGR of 26 percent to 2.3ZB by 2020.

We believe zettabytes need "zettastructure" – open source, software-defined, data-driven, hyperscale and autonomous infrastructure: true SDDC.

A MarketsandMarkets study estimates that the global market for SDDC will bring in $25.6bn in revenue in 2016, growing at a CAGR of 26.6 percent to $83.2bn by 2021. So SDDC is a real market, growing really fast.

But these figures pull together the components of the IT stack only: software-defined infrastructure (SDI), composed of software-defined networking, storage and computing. The study includes services, such as consulting, integration and deployment, but it only counts the IT infrastructure stack, north of the rack. It leaves out the south, or physical MEP infrastructure side.

In our opinion, mechanical and electrical infrastructure (thermal and power management) systems must also become software-defined, where the software is data-driven, predictively analytical, policy-based and tightly integrated into IT-stack performance management.

> *The true jump to SDDC only occurs when automation is brought to bear on the physical infrastructure*

The true jump to SDDC occurs only when software automation and data-driven analytical intelligence are brought to bear on the physical critical-environment infrastructure.

Critical environment functions have been handled under the catch-all category of data center infrastructure management (DCIM) and, more recently, a higher order of function known as data center service optimization (DCSO), which seeks to integrate DCIM with IT service management (ITSM). However it is done, we need to see an end to the old silos.

For years in the industrial world, IT and operational technology (OT) have been treated as separate disciplines. Now fresh thinking and new technologies are giving IT the ability to automate and use OT data.

There are those who don't think we need to integrate and software-define the full stack. *DCD* disagrees. During the past decade we have learned the folly of treating the logical and physical sides of the critical-environment world as different countries (if not planets).

When design and engineering professionals on both sides of the border speak two different languages, this creates threats to uptime, availability, resiliency and efficient IT performance.

In the early days of virtualization, the pace of change was hard for facilities engineers to keep up with, as server, storage and networking technology advanced with every server refresh. Power and cooling were static for the life of the facility – at least a decade. That is no longer true.

For now, the true SDDC may be limited to those organizations with deep pockets and monolithic applications – the vertically integrated hyperscalers and cloud services providers that can push out the boundaries of data center-as-a-service. But anyone requiring DevOps-style Internet-facing agility at the application and workload level will increasingly want these characteristics from its in-house or outsourced data center-as-a-service provider. To meet the demands placed on them, data centers must become open source, full-stack integrated, software-defined and autonomous, right down to the lowest level of their infrastructure.

We must move towards a world where cloud services and capacity will not require human touch, except where humans are unwilling to let control go to software.

"Software is eating the world," Marc Andreessen famously opined in the ancient history of 2011. That is now coming true in ways he could not have predicted. ●

DCD > Webinars | 2017

Juniper NETWORKS

NEW – Save the date!

The Path to the Cloud and the Role of the Network

Thursday 11 May
15:00 BST

Speakers confirmed:

Andy Ingram
Worldwide Data Center Solutions Lead
Juniper Networks

Stephen Worn
CTO
DCD

Everything is moving to the cloud. Or is it? In fact, many applications are far from cloud ready, and some never will be. But almost every new application will be cloud-enabled. How do you evolve your data center into a private cloud and drive Hybrid IT by connecting to public clouds, and still cater for the needs of your legacy software?

Attend this webinar and take away practical ideas to enhance your cloud and data center planning projects.

This webinar will examine:

- Why many applications don't necessarily move easily into the cloud.
- The difference between older "mode 1" applications that lack flexibility and modern "mode 2" applications built specifically for the cloud.
- What is needed in order to have an architecturally coherent network environment that can support these multiple generations of applications.
- The importance of topology in the cloud-ready data center network, and the various choices available.
- How to use an open architecture ecosystem to drive automation by providing software control of the network and integrating it into the virtualization capabilities of the data center.

▶ **Register:** http://bit.ly/2o4NCCK

For more information please contact **webinars@datacenterdynamics.com**