



Issue 58 • September 2025
datacenterdynamics.com

THE UK'S COPPER SWITCH

OFF

**HOW BT & OPENREACH
ARE APPROACHING THE
END OF COPPER**

ALSO INSIDE:
JOHOR IN FOCUS
BUILDING A SUPERCOMPUTER
MOVIE STUDIOS IN THE CLOUD



**Reliable, Scalable,
and Efficient**

Data Center Cooling Solutions

In mission-critical data centers, maintaining precise cooling is essential for uptime, performance and efficiency. The Belimo Energy Valve™ ensures consistent differential pressure in direct-to-chip cooling applications, optimizing heat transfer and preventing thermal fluctuations that can impact server performance.

Belimo solutions provide:

- Reliable temperature regulation to prevent IT equipment failures
- Scalable solutions to support increasing server loads
- Efficient energy management to reduce waste and operational costs

Experience Maximum Uptime and Optimal Performance with Belimo.



Learn More
www.belimo.com/datacenters



50
YEARS

**Together
to the Top.**

BELIMO®

Contents

September 2025



6 News

17 The UK's copper switch-off

Inside the BT tower, and the telco's plan to switch the UK's network from copper to fiber

31 The Critical Power supplement

Powering the next generation of data centers

47 Drilling for data

Can geothermal power match the grand ambitions of the hyperscalers?

52 Going local in Europe

Data center veteran Adriaan Oosthoek talks about his new firm Portus and the rise of Europe's regional Edge

54 The Arc of Trout

Jim Trout is CEO of a digital infrastructure veteran now at the helm of CleanArc Data Centers, a new company matching sites with renewable energy



57 Johor in focus

Malaysia's fastest growing data center market is benefiting from its proximity to Singapore

63 The big fightback

Vandals have been targeting telecoms infrastructure with increasing regularity in recent years, and the networks have had enough

69 How to build a supercomputer

DCD gets up close and personal with Hunter, the most powerful supercomputer at the Stuttgart HPC center

75 A story Untold

Untold Studios is a visual effects studio based entirely in the cloud. It reflects a trend in Hollywood towards using virtual infrastructure

81 The coming storm

Space weather could pose an existential threat to the satellite industry unless its impact can be better understood

86 The heat is on

The UK lags other countries in Europe when it comes to connecting district heating systems to data centers. But that could be set to change

94 The US fiber frenzy

The biggest names in the US telco market are all busy boosting their fiber networks through mergers and acquisitions

98 Data center security in the AI girlfriend era

A consequence of the AI boom you may not have considered before



Because it's **your assets** on the line.



*Mission Critical Computer Room
Air Handler (MCH CRAH)
Up to 90kW*



*Mission Critical Vertical
Computer Room Air Handler
(MCV CRAH)*



*YVAM Air-cooled Magnetic Bearing
Centrifugal Chiller*

From chillers optimized for low-GWP refrigerants to purpose-built air handling solutions, **YORK® is shaping the next era of high-density data center thermal management.** Discover how customer-driven innovation, global scalability and engineered consistency make YORK® the powerhouse behind data center solutions from Johnson Controls that are more efficient and more intelligent than ever before.



**Thermal
Management**



**Building
Controls**



**Facility
Security**



**Fire
Protection**



**Proactive
Services**



Learn more about how YORK® is
engineering the future of data centers:

[YORK.com/data-centers](https://www.york.com/data-centers)



From the Editor

The (other) Tower of London

If you ask someone to draw the London skyline, among images of St Paul's Cathedral, the London Eye, and the Cheesegrater, a telco tower stands supreme.

A journey up the tower, and down copper lane

For our cover feature (p15), we were one of the last journalists to visit the iconic BT Tower, to understand how BT and Openreach are preparing to move out of it and thousands of exchanges as they shift from copper to fiber. Plus, we profile the entire history of the company's copper usage.

The slow death of the UK's copper networks

The power beneath your feet

As AI data center demands increase, so is the industry's willingness to look beyond traditional power sources. Hyperscalers like Meta and Google are partnering with innovative new geothermal companies to unlock energy from the ground (p47).

Portus' head

A new company hopes to find space at the regional Edge (p52). We speak to Portus' chairman about the company's plans to build small-scale data centers in Europe, targeting the kinds of customers that have been forgotten in the AI boom.

The big fish

Another relatively new business hopes

to target larger customers. CleanArc CEO Jim Trout talks us through the company's plans to build hyperscale-focused data centers matched to renewable energy (p54).

Johor rising

With Singapore filling up, neighboring Johor has pitched itself as the natural successor. We profile Malaysia's plans to become the next big data center hotspot (p57).

Thieves and terrorists

Telecoms infrastructure is in the crosshairs. Thieves are stealing copper, while conspiracists are attacking towers. We look at the growing problem (p63).

Hunting for supercomputers

We travel to Germany to profile the Hunter supercomputer, featuring the latest AMD hardware (p69).

Telling untold tales on the big screen

The compute and networking requirements of visual effects are in focus in this profile of Untold Studios (p75).

Disaster in space

The satellite industry needs to do more to prepare for space weather events. We look into what can be done (p81).

Keeping the heat

We go behind the UK's challenges with district heating, and consider whether it could follow a Nordic model (p86).

Plus:

A critical power supplement, Imperial College's CIO, M&A in telecoms, AI girlfriends, and much more!

\$5.3 Billion

The price Telereal paid for more than 6,000 BT exchanges in 2001 (inflation adjusted)



Sebastian Moss
Editor-in-Chief

Publisher & Editor-in-Chief
Sebastian Moss

Managing Editor
Dan Swinhoe

Senior Editor
Matthew Gooding

Telecoms Editor
Paul 'Telco Dave' Lipscombe

Compute, Storage, & Networking Editor
Charlotte Trueman

Cloud & Hybrid Editor
Georgia Butler

Energy & Sustainability Senior Reporter
Zachary Skidmore

Junior Reporter
Jason Ma

Head of Partner Content
Claire Fletcher

Partner Content Manager
Farah Johnson-May

Copywriter
Erika Chaffey

Designer
Eleni Zevgaridou

Media Marketing
Stephen Scott

Group Commercial Director
Erica Baeta

Conference Director, Global
Rebecca Davison

Live Events
Gabriella Gillett-Perez
Tom Buckley
Audrey Pascual
Joshua Lloyd-Braiden

Channel Management Team Lead
Alex Dickins

Channel Manager
Kat Sullivan
Emma Brookes
Zoe Turner
Tam Pledger
James Raddings

Director of Marketing Services
Nina Bernard

CEO
Dan Loosemore

© 2025 Data Centre Dynamics Limited All rights reserved. No part of this publication may be reproduced or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, or be stored in any retrieval system of any nature, without prior written permission of Data Centre Dynamics Limited. Applications for written permission should be directed to the editorial team at editorial@datacenterdynamics.com. Any views or opinions expressed do not necessarily represent the views or opinions of Data Centre Dynamics Limited or its affiliates. Disclaimer of liability: Whilst every effort has been made to ensure the quality and accuracy of the information contained in this publication at the time of going to press, Data Centre Dynamics Limited and its affiliates assume no responsibility as to the accuracy or completeness of and, to the extent permitted by law, shall not be liable for any errors or omissions or any loss, damage or expense incurred by reliance on information or any statement contained in this publication. Advertisers are solely responsible for the content of the advertising material which they submit to us and for ensuring that the material complies with applicable laws. Data Centre Dynamics Limited and its affiliates are not responsible for any error, omission or material. Inclusion of any advertisement is not intended to endorse any views expressed, nor products or services offered, nor the organisations sponsoring the advertisement.

Dive even deeper

Follow the story and find out more about DCD products that can further expand your knowledge. Each product is represented with a different icon and color, shown below.



Events



Intelligence



Debates



Training




Awards



CEEDA

News

Meta goes gigascale with Prometheus and Hyperion



Meta to invest “hundreds of billions of dollars into compute to build superintelligence,” with several multi-GW data center clusters.

CEO Mark Zuckerberg in July said he has aggressively ramped up the company's AI effort, spending hundreds of millions on key hires and billions on acquisitions, as the company looks to leapfrog competition.

“We’re also going to invest hundreds of billions of dollars into compute to build superintelligence. We have the capital from our business to do this,” he said.

“We’re actually building several multi-GW clusters. We’re calling the first one Prometheus and it’s coming online in ‘26.

“We’re also building Hyperion, which will be able to scale up to 5GW over several years. We’re building multiple more titan clusters as well. Just one of these covers a significant part of the footprint of Manhattan.”

The Ohio Prometheus development will use gas turbines and is based on a new data center design focused on speed to deployment, *SemiAnalysis* reports.

Meta is also developing a massive \$10 billion data center in Richland Parish, northeast Louisiana, known as Hyperion. First announced last year as a four million sq ft campus, it is expected to take until 2030 to be fully built out.

By the end of 2027, it could have as much as 1.5GW of IT power.

This quarter saw Meta detail plans to sell some of its in-development data center portfolio for co-development.

“At the beginning of June 2025, we approved a plan to dispose of certain data center assets and reclassified \$2.04 billion of these assets as held-for-sale,” the company said in the SEC filings for its latest quarterly earnings.

“These assets mostly consisted of construction in progress and land, which are reported at the lower of their carrying amounts or fair values less costs to sell.”

Meta said these assets are expected to be disposed of within the next 12 months through a contribution to a third party for the purpose of co-developing data centers.

Details on which data centers the company might sell – as well as potential buyers – haven’t been shared. The fact that the sites are under construction suggests the company will likely lease them back once sold.

Meta has around 30 data center campuses in operation or development globally. The majority of its footprint is within the US. It is also a major leaser of facilities. The first facilities using a new design that caters for GPUs and liquid cooling are due live next year.

NEWS IN BRIEF

Meta using tents as temporary data centers

Social media giant Meta is using tents to house temporary data center capacity. CEO Mark Zuckerberg said: “We’re basically building these weatherproof tents... and building GPU clusters inside them.”

NREL publishes US data center infrastructure map

The US National Renewable Energy Laboratory (NREL) has released a map illustrating the US data center sector and its relationship with the power and fiber infrastructure that serves it.

Construction worker dies at DC site in London

A construction worker died after an incident at the site of a Pure Data Centers development in London, UK. The incident, which took place in Brent Cross, is being investigated.

Global layoffs at Oracle


Oracle is in the midst of laying off a large number of workers around the world. The company hasn’t revealed the full scope, but have impacted Oracle Cloud teams across the US, India, the Philippines, and Canada. The move follows the layoff of several hundred OCI employees last year. Staff from Cerner are said to be impacted.

Figma spends \$\$\$ with AWS

Design tool company Figma has revealed its daily cloud computing spend. In its initial public offering filing, the company details that it is spending a massive \$300,000 on cloud computing services daily, with AWS as its primary cloud platform. In May it signed a five year deal committing to a minimum spend of \$545 million.

Hyperscalers prepare for 1MW racks

Google has joined Meta and Microsoft’s collaboration project on a power rack the companies hope will help them reach rack densities of 1MW. The companies are working on a new power rack side pod called Mount Diablo that will enable companies to put more compute into a single rack. Specs will be released to OCP this year.



US utilities see data center connection requests grow

Utilities across the US have seen connection requests from data center customers reach new highs.

PJM Interconnection, the largest Regional Transmission Organization (RTO) in the US, expects to see a peak load growth of 32GW from 2024 to 2030, with 30GW from data centers.

The utilities that will provide the power generation assets to serve this demand are seeing huge demand from data centers.

In August, First Energy said its long-term pipeline for data center load increased by more than 80 percent to 11.1GW from 6.1GW in February, with its concentrated data center load through 2029 increasing by around 25 percent to 2.7GW. First Energy utility companies serve more than six million customers across Ohio,

Pennsylvania, New Jersey, West Virginia, Maryland, and New York.

New Jersey's largest electrical and gas utility, Public Service Enterprise Group Incorporated (PSEG), has seen large-load interconnection requests rise to 9.4GW, driven almost entirely by the data center sector. In total, the utility said that large load requests had grown 47 percent since Q1. Despite the pipeline growth, LaRosa said that he expects only between ten to 20 percent of those requests to come online.

PPL Electric, which primarily serves Pennsylvania, has seen its pipeline of data centers in the "advanced-stage" of interconnection grow to 14GW, a 33 percent jump in the three months to August. It means PPL's data center load could swell from 800MW in 2026 to 14.4GW in 2034,

with some 5GW already publicly announced.

In California, Pacific Gas and Electric Company (PG&E) has reported another increase in its data center pipeline in its latest report.

The utility recently said that it now has 10GW of capacity in its pipeline, expected for connection over the next ten years. This is a rise from May, when it reported an 8.7GW pipeline, and February, when the company said it had a pipeline of 5.5GW.

Of the 10GW, 17 data center projects with a combined capacity of 1.5GW are in the final engineering phases, with operations slated to begin between 2026 and 2030. The majority are in San Jose, Silicon Valley, and the greater San Francisco Bay Area, but also include some in the Central Valley and Sacramento.

Meanwhile, Detroit-based utility DTE Energy is in talks with data centers representing up to 7GW of potential new capacity around Michigan.

"We are in advanced discussions with multiple hyperscalers for over 3GW of new load, and are having ongoing discussions with multiple other data center operators for an additional 4GW," DTE president and CEO Joi Harris said during the company's Q2 earnings call in July.

American Electric Power (AEP), which serves 11 states including Indiana, Kentucky, Louisiana, Michigan, Ohio, Oklahoma, Tennessee, Texas, Virginia, and West Virginia, reported a pipeline of 18GW of incremental data center load growth to 2029.

DOE offers federal land to DCs



The US Department of Energy (DOE) has announced the first four sites for AI data center developments on Federal land.

The department said that it invites private sector partners to develop at Idaho National Laboratory, Oak Ridge Reservation, Paducah Gaseous Diffusion Plant, and Savannah River Site. Developments are expected to include both data centers and related power infrastructure.

"By leveraging DOE land assets for the deployment of AI and energy infrastructure, we are taking a bold step to accelerate the next Manhattan Project - ensuring US AI and energy leadership," said Energy Secretary Chris Wright.


"These sites are uniquely positioned to host data centers as well as power generation to bolster grid reliability, strengthen our national security, and reduce energy costs."

The move to open up Federal land for the data center sector was first pushed by the Biden administration and furthered by the Trump administration.

In April, the DOE said that it had 16 potential locations, including the four sites detailed here, mostly at national laboratory sites. The DOE is still evaluating additional sites for development.

President Trump signed an executive order removing clean energy stipulations from the data center proposal and instructed agencies to streamline environmental reviews and permitting.

AT&T to snap up \$23bn spectrum licenses from EchoStar, Boost plans to become an MVNO



AT&T has announced that it will acquire certain wireless spectrum licenses from EchoStar for approximately \$23 billion.

As part of the deal, the carrier and EchoStar have agreed to expand their long-term wholesale network services agreement, enabling EchoStar to operate as a hybrid mobile network operator as it provides wireless services under its Boost Mobile brand.

The nature of the deal sounds very much like EchoStar will become an MVNO instead of the fourth carrier it hoped to become to compete in the market against the "Big 3" telcos - AT&T, Verizon, and T-Mobile.

AT&T noted that it will be the primary network services partner to EchoStar as it continues to serve wireless customers.

The deal will see AT&T acquire approximately 30 MHz of nationwide 3.45 GHz mid-band spectrum and approximately 20 MHz of nationwide 600 MHz low-band spectrum for approximately \$23bn in an all-cash transaction.

According to AT&T, these licenses "cover virtually every market across the US" - more than 400 markets in total. The carrier adds that the acquisition significantly strengthens AT&T's low-band and mid-band spectrum holdings.

"This acquisition bolsters and expands our spectrum portfolio while enhancing customers' 5G wireless and

home Internet experience in even more markets," said John Stankey, chairman and CEO, AT&T.

AT&T said it wants to deploy the mid-band licenses "as soon as possible."

The transaction is expected to close midway through next year, subject to certain closing conditions, including regulatory approvals.

Beyond pushing AT&T's 5G coverage, the carrier adds that the spectrum will bolster its push to move customers from legacy copper-based phone and Internet services to next-generation connectivity such as AT&T Internet Air and AT&T Phone - Advanced in areas that AT&T will not reach with fiber.

In other US telecom news, T-Mobile closed its \$4.4 billion acquisition of UScellular's wireless operations, including the company's 4.5 million customers.

The announcement of the deal came August 1 after receiving approval from the Federal Communications Commission (FCC) in September.

Post-acquisition, UScellular is rebranding itself as Array Digital Infrastructure, focused on the 4,400 cell towers the company still owns.

This quarter also saw T-Mobile acquire Minnetonka-based U.S. Internet (USI) as part of its fiber expansion plans.

The firm operates a fiber network in Minneapolis.

Pennsylvania becomes a gigawatt-scale market

Pennsylvania has become the latest US state to see a ballooning in proposed data center campuses.

Data center analyst firm DC Byte reports the Keystone State has seen its total IT load in planning and operation jump from 231MW in 2021 to 7.8GW today. This figure includes projects that are live, under construction, committed, and early-stage.

Most of Pennsylvania's colocation market is driven by wholesale colocation, which has jumped from 2MW in early-stage plans and 4MW committed in 2021 to around 3GW early-stage and 2GW committed in 2025. Pennsylvania's hyperscale market has also seen growth, going from no early-stage projects in 2021 to 616MW worth of them by 2025, but the market size is much smaller.

In June, Amazon announced plans to invest \$20 billion in data centers in the state. The company is developing adjacent to the Susquehanna nuclear plant.

In July, data center and energy firms, including Blackstone and Energy Capital Partners, pledged to invest \$90bn in Pennsylvania's digital infrastructure.

July also saw CoreWeave announce plans to invest \$6 billion in an initial 100MW of capacity in Lancaster, with potential to reach 300MW, co-developed by Chirisa Technology Parks and Machine Investment Group, with backing from Blue Owl Capital.

Scotland sees influx of data center proposals filed

Scotland might finally be set to fulfill its data center potential.

Blessed with abundant renewable power and available land, the country has long had a tiny data center market with a handful of local players.

Recent months, however, have seen multiple proposals for large-scale project filed, led by local renewable energy developer Apatura.

The firm has filed at least four applications for data center projects across the country, including Falkirk, Motherwell, Cumnock, and the Scottish Borders, totaling more than 1GW of capacity.

Founded in 2014, Apatura has battery energy storage projects planned across Scotland; the company says it has 10.6GW worth of grid-secured energy projects, including 1.8GW of grid connections targeting data center developments.

This quarter also saw Shelborn Asset Management file to turn a former bank office site in Edinburgh into a data center. The company had permission for a new 800,000 sq ft office complex, but has pivoted instead to data centers.

The STULZ logo consists of the word "STULZ" in a bold, white, sans-serif font, centered within a red rectangular background.

CLIMATE. CUSTOMIZED.



Precision Cooling

3 kW - 1080 kW of Cooling Capacity

STULZ designs, manufacturers, and supports precision cooling and humidity control solutions for mission critical applications.



User Driven | Custom Designed | Purpose Built

www.stulz-usa.com

Fire at Telecom Egypt data center kills four, 26 hurt

Four people died after a fire engulfed a Telecom Egypt data center in Cairo, Egypt.

The fire, which happened July 7, occurred at the 10-story facility and lasted for hours.

In addition to the four people who died, 26 people were injured.

The cause of the fire is not yet known, though the governor's office confirmed that it started on the seventh floor of the building in the Ramses area of Cairo.

The fire caused temporary mobile network and Internet outages. It also impacted other services such as online payments and air traffic control.

"A fire broke out this evening in one of the equipment rooms at the Ramses switchboard of the Telecom Egypt company, which led to a temporary disruption of telecommunications services," said the National Telecom Regulatory Authority (NTRA) in a statement.

During the inquiry, Egypt's Minister of Communications and Information Technology, Amr Talaat, said the fire suppression system in the building was overwhelmed by the intensity of the flames.

Telecom Egypt posted a statement expressing its "deepest condolences" to the people who died in the fire, saying it was supporting the families of the deceased and injured.

OpenAI expands into Europe, plots Norwegian site



Generative AI firm OpenAI has expanded its Stargate data center project internationally again, targeting a facility in Norway.

Announced in July, the company will work with data center company Nscale and engineering firm Aker for an initial 230MW facility. The site could add another 290MW.

The data center will target to deliver 100,000 Nvidia GPUs by the end of 2026, with plans to expand "significantly" in the years ahead.

Located in Kvandal, outside Narvik, northern Norway, the site was announced by Aker earlier the same month. Nscale also operates a data center in Glomfjord, southern Norway. The two companies will own the new site in a 50/50 joint venture.

OpenAI's Stargate project is a \$500bn effort to build massive data centers - originally across the US, but now globally - for the AI developer.

The likes of Oracle, SoftBank, and Abu Dhabi's MGX are named investors in the venture.

Crusoe is developing a large campus for Stargate in Texas. The Abilene campus, owned by Lancium, is expected to eventually have eight buildings and a total of 1.2GW of capacity.

OpenAI is also set to develop a campus in the United Arab Emirates.

The 1GW Stargate UAE site, to be built by G42 and operated by OpenAI and Oracle, will be located within a planned 5GW AI campus in Abu Dhabi. Set to be equipped with Nvidia's GB300 GPU systems, the first 200MW is set to go live sometime in 2026.

In the US, OpenAI is also reportedly exploring other Stargate data center options in states across the US including Arizona, California, Florida, Louisiana, Maryland, Nevada, New York, Ohio, Oregon, Pennsylvania, Utah, Texas, Virginia, Washington, Wisconsin, and West Virginia.

Beyond the US, previous reports have suggested that OpenAI was looking to develop up to ten Stargate data center projects globally.

Recent reports suggest the company is targeting India.

OpenAI's Stargate India project could reportedly total 1GW of IT capacity. Further details haven't been shared.


OpenAI recently said the company has increased its compute 15-fold since 2024. OpenAI CEO Sam Altman noted in August that generative AI could be a bubble, while at the same time promising trillions of dollars of data center and compute spend.

The company's CFO recently said that the company could one day offer cloud services.



Dan's Data Point

Two cranes collapsed at data center sites this quarter. July saw a crane at a Meta-owned development in South Carolina collapse, crushing a car. August saw another collapse at AirTrunk's site in Melbourne, Australia. Luckily, no one was hurt in either incident.



US govt buys stake in Intel

The US government has taken a 9.9 percent stake in semiconductor giant Intel.

In August, the US government took a 9.9 percent stake in the beleaguered chipmaker and designer for \$8.9 billion in Intel, with \$5.7bn coming from funds that would have been awarded through the CHIPS Act and \$3.2bn from funds from the Secure Enclave program.

"The United States paid nothing for these Shares, and the shares are

now valued at approximately \$11 billion Dollars," President Donald Trump said on Truth Social (the US did pay for the shares, but through already-allocated chip funds that Trump previously threatened to rescind).

"This is a great Deal for America and, also, a great Deal for INTEL."

The deal came a week after Intel CEO Lip-Bu Tan came under pressure from government officials over claimed links to China, with Trump calling for his resignation.

Home of Data

Tate.

Retro-Fit Ready

Retrofit Your Data Centre with Liquid Cooling Solutions from Tate



The compute demands of AI, GPUs, and accelerated workloads are outpacing your infrastructure cycle. Tate's liquid cooling manifolds make it easy to upgrade existing air-cooled data centres, allowing you to support increased densities without rethinking your architecture.

- Integrates into the hot aisle containment system
- Global production with local support
- Suitable for both new build and retrofit

AI Solutions, Fast Tracked.

tateeurope.com

Scan here
to learn more



Former Texas governor plans 11GW data center campus



A massive multi-gigawatt behind-the-meter campus is being planned in Texas close to the US government's main nuclear weapons development facility by former Texas governor and energy secretary Rick Perry.

In June, the Texas Tech University System (TTU System) announced a collaboration with recently founded energy and technology company Fermi America to develop a large AI and energy campus near Amarillo in the Texas Panhandle.

Located on approximately 5,800 acres, the Advanced Energy and Intelligence Campus will reportedly include "18 million square feet (1.67 million sqm) of data centers and generate up to 11GW of IT capacity from natural gas, solar, wind, and clean nuclear energy."

Fermi America, co-founded by Perry, will provide critical infrastructure energy solutions.

The company said geotechnical work has already begun on the campus and is expected to deliver 1GW of power by the end of 2026.

Fermi says its HyperGrid behind-the-meter campus will combine the "largest nuclear power complex in America, the nation's biggest combined-cycle natural gas project, utility grid power, solar power, and battery energy storage" into one development serving hyperscale data center companies.

The company said the planned campus is located at the "confluence of several of the nation's largest gas pipelines and located atop one of the nation's largest known natural gas fields."

Fermi also said the planned campus

is also in proximity to Pantex, the site is the US' primary nuclear weapons assembly and disassembly facility. It is located 20 miles (32 km) northeast of Amarillo, in Carson County.

"The Chinese are building 22 nuclear reactors today," added Perry. "We're behind, and it's all hands on deck. President Trump's first executive order spoke to the energy issue and how we must continue to make America energy dominant. No one does energy better than Texas, and Fermi America and the Texas Tech University System are answering the call."


James Richard "Rick" Perry was the 47th governor of Texas from 2000 to 2015 and secretary of energy under President Trump from 2017 to 2019. He unsuccessfully ran for the Republican presidential nomination in the 2012 and 2016 elections.

The campus plan also includes the construction of an academic and research facility featuring office, classroom, and conference spaces.

TTU System said the project will provide academic and research opportunities for faculty and students, including internships, employment, and workforce training and placement programs that will benefit all five TTU System institutions.

Fermi has since secured investment and financing from Macquarie, and announced partnerships with Westinghouse, Hyundai, and Doosan Enerbility.

The company has secured 600MW of gas turbines for the project.



Equinix signs three nuclear deals totaling more than 750MW

US colo firm Equinix has signed three major agreements with advanced nuclear developers that together could provide up to 774MWe of power.

The first deal was with Radiant, a microreactor company based in El Segundo, California. As part of the agreement, Equinix will purchase 20 Kaleidos microreactors, each designed to produce

1.2MWe and 3MWth.

The Kaleidos is a high-temperature gas-cooled reactor (HTGR) that uses TRISO fuel, a helium gas coolant, and prismatic graphite blocks. Radiant says its modular containerized system can be deployed to deliver on-site power within days.

The second agreement saw Equinix sign a Letter of Intent with ULC-Energy for a Power Purchase Agreement of up to 250MWe to power its data centers in the Netherlands. ULC-Energy, based in Amsterdam, is the exclusive Dutch development partner for Rolls-Royce SMR and selected the company's 470MWe light water small modular reactor as its preferred technology in 2022.

Equinix's third deal was a pre-order power agreement with Stellaria, a French nuclear company, for 500MWe to supply data centers across Europe. Founded in June 2023 by the French Alternative Energies and Atomic Energy Commission (CEA) and Schneider Electric, Stellaria is developing the world's first molten salt Breed & Burn reactor.


Last April, Equinix signed a pre-agreement with Oklo to procure up to 500MW of power from future fast fission plants. With the latest announcements, Equinix now has agreements covering more than 1GW of nuclear capacity.

Partnering for AI-Ready Data Centers

Transform your data center into a powerhouse of efficiency with Schneider Electric's AI-Ready solutions. Our end-to-end infrastructure services are designed to adapt to the growing demands of AI workloads, ensuring your operations remain resilient. Leverage advanced power and cooling systems tailored for high-density compute environments that maximize performance while minimizing environmental impact. Partner with us to redefine your data center's energy strategy and take the lead in the AI era.



Scan to Learn More



Life Is On

Schneider
Electric

Data4 agrees 40MW nuclear energy deal with EDF



European data center firm Data4 has signed a nuclear power supply agreement with electric utility EDF to power its French data center portfolio.

The Nuclear Production Allocation Contract (NPAC) will last for 12 years and see EDF supply 40MW of power from its operational fleet. As part of the agreement, Data4 will share a portion of the costs and risks associated with the volumes of energy actually produced.

The first delivery of power is expected in 2026, with EDF projected to supply an annual volume of around 230GWh.

"We are extremely proud to sign this first low-carbon CAPN with EDF, a long-standing partner and key player in the energy transition in France," said François Stérin, chief operating officer of Data4.

"This contract is not only a strong

environmental commitment; it is also a pillar of our growth and competitiveness strategy. Combined with power purchase agreements for renewable energies, it guarantees us reliable, resilient, and continuous access to low-carbon energy at a controlled cost over the long term, a considerable advantage for our customers and for the development of our infrastructure."

Data4 will be the first data center operator in France to sign such an agreement.

The NPAC contract structure was introduced this year, with EDF offering it via a European-wide auction mechanism, targeting large energy consumers (more than 7GWh per year) and electricity suppliers/producers with the capacity to physically receive electricity in France.

"EDF is delighted to have signed a nuclear generation allocation contract with Data4, a European data center operator. This partnership will enable Data4 to benefit from a low-carbon, competitive, and available electricity supply, which is a key asset for digital players. This approach is fully in line with EDF's commitment to the development of data centers in France as well as to the country's energy and industrial sovereignty," said Marc Benayoun, EDF group executive director in charge of customers, services, and territories.

The agreement builds on a ten-year Power Purchase Agreement signed last year between Data4 and EDF-owned Italian utility Edison Energia for 500GWh of energy from an Italian solar project.

The nuclear supply agreement is one of the first of its kind in the European market, with the majority of major deals involving traditional nuclear operators, primarily in the US market. Meta, Microsoft, and AWS have all signed long-term PPAs to offtake power from large-scale nuclear plants in the US.

The European market has not been as active. However, last month saw colo data center developer Equinix sign deals with two European advanced nuclear companies, agreeing to offtake 250MW from ULC-Energy in the Netherlands, and 500MW from Stellaria in France. Following the deal with Equinix, Stellaria announced the successful raise of €22 million (\$25.6m) as part of its latest funding round. EDF previously signaled its willingness to allow data center firms to colocate at its power plants in France.



Quantinuum valued at \$10bn

Quantum computing firm Quantinuum has received a pre-money valuation of \$10 billion following a \$600 million equity raise, double what the company was valued at after the close of its last funding round in January 2024.

The funding round, announced in early September, saw participation from Nvidia's venture arm NVentures and Quanta

Computer, in addition to previous investors, including JPMC, Mitsui, Amgen, and Honeywell.

Quantinuum said the funding would help the company to further scale its quantum computing technology in the lead-up to the launch of its Helios system later this year. According to the company, Helios will empower users "to solve classically intractable problems and explore the hybrid classical-quantum application space."

The system will be available as a Hardware-as-a-Service (HaaS) offering, with cloud access also set to be offered.

Quantinuum was founded in 2021 when Honeywell spun out its Quantum Solutions division and merged it with UK quantum computing startup Cambridge Quantum Computing. Honeywell owns a 54 percent stake in Quantinuum. IBM is also an investor.

In addition to the fundraising, Quantinuum announced it had been made a founding collaborator at Nvidia's Accelerated Quantum Research Center. The company has also recently partnered with Riken, SoftBank, Infineon, and the UK's Science and Technology Facilities Council Hartree Center, and has deployed infrastructure or provided access to its quantum computers in New Mexico, Qatar, and Singapore.



Introducing the Munters LCX: Hyperscale by design

Designed for liquid-cooled servers, our new LCX liquid-to-liquid Coolant Distribution Unit is built for efficiency, reliability, and ease of serviceability. It provides precise control of the technology fluid supply temperature and external pressure or flow, suitable for rejecting heat from cold plates, in-rack CDUs, and other liquid heat rejection devices.

What sets Munters apart is our ability to customize components, such as brazed plate heat exchangers and pumps, to optimize performance for each installation. Available in sizes from 500 kW to 1.5 MW, the LCX ensures efficient, scalable, and serviceable liquid cooling, tailored to the demands of modern high-density data centers.



Learn more at munters.com



Learn more about
Munters LCX



Data center networking
is a playoff game.
And we're bringing our game face.



Score your biggest wins ever with Nokia as your data center networking partner. We have all the connections to win in the AI era:


- Fabrics that enhance connectivity within the data center
- IP and optical technologies that **seamlessly interconnect** data centers, clouds, WAN and the Internet
- Network management **and** automation powered by AI, **for AI**



Let's team up today

NOKIA

The end of the line: The UK's copper switch-off



Dan Swinhoe
Managing Editor

How BT & Openreach are approaching the end of copper

On the banks of London's River Thames is a school with a football pitch (soccer to those of an American persuasion) – one of the few such sports grounds within the old City of London proper. Those attacking the goal on the far side of the pitch will be greeted with a great view of St. Paul's Cathedral.

But those with an interest in digital infrastructure might be more interested

in the two buildings seen in front of that iconic church. In the middle-ground is British Telecom's (BT) Faraday House, part of what was one of the first telephone exchanges in London and a notable piece of telecoms history.

And immediately behind the opposition goal is the telco's Baynard House, a brutalist monolith and probably the only data center with a wall regularly used to stop balls being kicked out of bounds.

Both of these buildings, and thousands more like them across the UK, make up the core of the UK's copper network. But not for much longer.

Though the first next-generation fiber networks date back to the 1970s, copper has remained the dominant way to deliver phone and Internet service well into the new millennium.

But as today's fiber networks reach

more and more of the furthest reaches of the UK and countries all over the world, time is now being called on legacy copper networks en masse.

Copper networks and telephone exchanges

The Public Switched Telephone Network (PSTN) is the traditional analog network in the UK that uses copper cables to carry landline voice calls between users. Beyond phone calls, the copper network also supports ADSL broadband services as well as the likes of alarm systems, medical devices, and monitoring equipment.

(sometimes known as Central Offices or Wire Centers), where the physical phone lines are interconnected and signals redirected between properties. Once, these would have had human switchboard operators manually connecting calls, before the process became automated from 1912 onwards.


Openreach showed *DCD* around the BT Tower, the company's iconic inner-city radio tower and telephone exchange, ahead of its planned closure and redevelopment into a hotel.

The company showed us PSTN rooms filled with System X and System Y/AXE10 switching gear, the MDF, its customer

into the room above, where they meet the main distribution frame (MDF), a large unit where the external copper network meets the internal networks – both Openreach's and its communication provider customers.

These rooms don't feature huge amounts of cooling equipment and the termination blocks have no blinking lights or fans whirring away. The room *DCD* saw feels more like an old garage than a key interconnection point.

More reminiscent of traditional data center white space is the multi-user area (MUA), where Openreach's communication provider customers can



"It has always started with 'this is why you can't do it,' but we've crossed the rubicon into 'how are we actually going to do that?'"

*>>Raj Chadha
Openreach*

The UK's first telegraph networks launched in the 1830s, with the first central telegraph station opening in London in 1849, and were nationalized by the Post Office around 1870. The first copper phone lines and public telephone exchange in Europe opened in London in August 1879. Launched by the Telephone Company (Bells Patents) Ltd at 36 Coleman Street, it boasted around a dozen subscribers and relied on overhead wires. It, and the 1,565 others that had popped up in the intervening years, were nationalized in 1912 by the General Post Office, giving the government a near-total monopoly of the country's nascent phone network.

Copper networks come with large numbers of exchange buildings

equipment hosting area, and the more modern fiber equipment that is all to be removed ahead of the company's exit from the site.

Even today, these exchanges don't feel like your average data center. While they generally have a cable chamber or vault of the kind found in many of today's facilities, these are typically larger and filled with thousands of strands of copper, some dating back decades, as well as more modern fiber wires. Old paper-insulated copper networks – which are still in operation in some places – even need their own systems pumping compressed air down the line to ensure no moisture can get in anywhere there's a tear in the lining.

From the chamber, the wires travel

place racks of networking gear that can handle both copper and fiber services.

Every house with a phone had at least one physical line that was wired to a nearby telephone exchange. Most people in the US or Europe will live near one and probably never even realize it. There are thousands in operation across the UK, and major cities would have multiple exchanges to handle the large number of phone lines – a website documenting the UK's exchanges notes there have been more than 100 across Greater London alone.

From the outside, exchanges vary in style. Coming from a different age, many are more aesthetically pleasing than your standard grey box modern data center, and were designed to fit in with



their surroundings. Due to the need for proximity to end users, exchanges are often located in the center of towns and cities, rather than on the outskirts.

Exchanges range from divisive brutalist structures, boring nondescript offices, and glorified sheds, to grand neo-Georgian beauties reminiscent of the lush art deco telephone exchanges of New York – many of which still serve as major interconnection hubs today. Openreach notes the UK's smallest exchange is located on the Shetland Isle of Papa Stour, serving just 14 homes; the largest still in operation, in the Oldham area of Manchester, serves more than 45,000 properties.

The UK's copper network and exchanges were built over a period of decades by incumbent operator BT. The telco was long government-owned and part of the Post Office until being re-privatized in the mid-1980s. Today, the copper network and exchanges are operated by Openreach, BT's wholesale network arm that offers fixed wireless access on a wholesale basis to communications providers across the UK.

But after more than 100 years, the telco has plans to shut down its copper network, and close thousands of exchanges across the UK.

This mammoth undertaking will have a massive impact not only on Openreach,

but also its hundreds of wholesale customers and, in turn, their millions of end users.

A shutdown for the ages

Copper's days have been numbered since 2017, when BT officially announced plans to shut down its PSTN and ISDN networks by the end of 2025 – a date later shifted back to December 2027 – and transition to a single all-digital Internet Protocol (IP) fiber network.

The telco is swapping out PSTN, ISDN (Integrated Services Digital Network – which allows voice and data to be sent

over the copper PSTN wires), and ADSL with full fiber-to-the-premises (FTTP) offerings and the slower single order generic Ethernet access (aka SOGEA, which offers fiber-to-the-cabinet and then copper from cabinet to an individual property).

Closing the thousands of telephone exchanges that the copper lines come into is a longer process that will happen in phases from this year until well into the 2030s. While tied to the PSTN shutdown, the exchange exit is technically a separate program; with the PSTN turned off, many exchanges will be largely idle and not be worth the effort and money to operate.

"We want everyone off copper, because copper has to terminate at a local exchange, and that's the fundamental problem here with closing exchanges," says Raj Chadha, senior manager, migration strategy, commercial and programme, at Openreach. "There is no case to build fiber everywhere and keep copper there at the same time."

He adds that the exit program is driven by redundancy: "Fiber does not need that many exchanges, and does not mean that many exchanges in London."

Openreach, which maintains BT's telephone cables, ducts, cabinets, and exchanges, was formed in 2006 to ensure rival telecom operators had equal access to BT's local network. It doesn't deal with



Credit: BT Group



HISTORY EXCHANGE: THE BT TOWER

Some exchanges exist quietly in sleepy corners of sleepy villages. Others loom over the landscape. The most notable is probably the BT Tower in central London.

What was then the Post Office Tower was launched in 1965 by Prime Minister Harold Wilson. At 620 feet (189 meters), it was Britain's tallest building at the time. It cost around £2.5 million to build – with the bill rising to £9 million including all the equipment (around £151.9m/\$206m in today's money).

The tower was designed to carry airdials for the Post Office microwave network covering some 130 stations throughout the country including what was then the Post Office satellite earth station at Goonhilly. It was built to host microwave antennae and replaced a shorter steel tower built on the roof of the neighboring Museum telephone exchange in the late 1940s to provide a television link between London and Birmingham.

A system of broadband and microwave radio links was proposed to meet the expected growth of the long distance telephone and television networks in 1954, according to BT. Two options were considered: a ring of radio stations on the outskirts with broadband links extended to central London by cable, or a single large radio station near the center. The latter idea won out.

The tower and the four-story building at its base were equipped to handle 150,000 simultaneous telephone connections and provide 40 channels for black and white or colour television.

A Grade II-listed building, the tower was open to the public until 1971 and featured a revolving restaurant operated by Butlins. It revolved once every 23 minutes, but was closed to the public after a bomb exploded in the roof of the men's toilets. Luckily

no-one was injured in the blast. The Angry Brigade, a far-left anarchist collective, claimed responsibility, while a call was also made by a person claiming to be the Kilburn Battalion of the IRA.


The restaurant remained opened by invitation until 1980, when Butlins' lease expired. BT and Openreach still use the tower's observation deck for corporate events. *DCD* was told we are likely one of, if not the last, journalists to go up the tower before it closes.

Despite its obvious nature and its public restaurant, the tower was, for a period, designated a classified location under the Official Secrets Act.

Hotel operator MCR announced the acquisition of the BT Tower from the telco in early 2024 for £275 million (\$346.8m). The hotel chain plans to turn the tower into an "iconic" hotel. MCR said BT will take a "number of years" to fully vacate the premises due to the scale and complexity of the work to move technical equipment.

MCR has 150 hotels across 37 US states, and this is its first project outside the US. It is aiming to bring the rotating restaurant back into public use.

Full details haven't been shared, but according to an interview with the *Sunday Times*, "billions of miles of telephone wire" from inside the tower will be woven into hotel wallpaper in a nod to the site's roots. Rotary phones in rooms are reportedly being considered. An infinity swimming pool is also said to be on the cards for the hotel, while Openreach tells *DCD* the exchange's main distribution frame is also set to be kept and made into a feature. ■



"We don't have control of providing into the end customer, and that becomes really difficult if you're trying to close an exchange. If the CP hasn't done enough, how do you manage that?"

*>>Raj Chadha
Openreach*

the end-users of its network, with its customers being the communications providers (CPs) that serve other businesses and/or the general public (including BT). It serves more than 500 ISPs and communications firms across the UK, and is leading the copper network and exchange shutdown.

Openreach currently operates some 5,600 telephone exchanges across the UK; most of which are for copper and other legacy services, with the company running fiber services from around 1,000 exchanges, or Openreach Handover Points (OHPs).

Customers using the firm's wholesale line rental and Local Loop Unbundling services, both of which are copper based, will need to move clients elsewhere, ideally to fiber offerings from remaining OHPs. Between the 2027 PSTN switch off and exchange exits, Openreach will continue to provide copper services from closing exchanges via interim metallic path facilities - copper-based services for providing broadband and voice services - and SOTAP, a copper broadband and IP voice service to replace ADSL in areas where fiber alternatives aren't available.

Consultations on closing of the 4,600 non-OHP sites started back in 2020, with BT saying it aimed to end leases at around two-thirds of exchanges leased from landlord Telereal Trillium. Even then, Openreach said it would be a "complex, long-term issue with implications for the whole industry."

In 2023, Openreach revealed plans to close 103 legacy exchanges by December 2030, starting with a trial of five sites, later reduced to three: Deddington, Oxfordshire; Ballyclare, Northern Ireland; and Kenton Road, London. All three are due to be fully decommissioned by the end of May 2026, with Deddington set to be the first to go. The next 105 exchanges (and 108 total) of the initial tranche are due to close in phases from 2028 to 2031. The remaining 4,500 will be closed over the course of the early 2030s.

A number of OHP exchanges - around 17 of the initial sites and a further 20 in the long run - will also be closed and consolidated into other OHP sites, leaving the final total of fiber-serving exchanges at around 960 facilities.

This isn't the first time the BT group has shuttered exchanges - the company has closed more than a handful of sites in the last 100 years or so - but the speed and scale of the current transition is unprecedented. The fact that migrating services need to be moved to fiber, instead of a like-for-like copper network, is another complicating factor.

Uptime and sustainability

According to BT, the PSTN can "no longer sustain modern requirements and is becoming increasingly fragile and prone to failure," and stats do suggest the network is becoming more brittle with age.

Ofcom's 2024 Connected Nations

report showed the number of significant PSTN resilience incidents reported increased 45 percent on the year prior - though the total number of lost hours declined as there were fewer customers on copper to impact.

The report shows there were 1,523 significant resilience incidents in 2024, up from 1,209 in 2023 and 1,281 in 2022. Hardware faults accounted for under a third of total lost hours; Ofcom noted PSTN incidents were growing due to equipment being beyond its intended lifespan and the reduction of qualified personnel within industry with experience of these legacy technologies.

BT also has a 2030 carbon-neutral goal it will be hard-pressed to meet if it doesn't turn off its copper services.

"We can't get a net-neutral industry without removing the PSTN and copper. It just has to happen," says Chadha.

European telecoms lobbying group Connect Europe has suggested a fiber-to-the-home line improves energy efficiency by at least 80 percent compared to copper, requiring about 90 percent less power and halving an operator's opex costs. Norway's Telenor has said it expects to save up to 100GWh annually after its copper switch-off - around an eighth of its total energy consumption across its fixed and mobile networks. A Telefónica report suggested the electricity consumption of fiber equipment is 10 times lower per megabyte emitted than



copper equipment.

Though it didn't break out network splits, BT has previously said it alone consumed nearly one percent of the UK's electricity annually as recently as 2022; it purchased some 2.5TWh of electricity in 2021. It has also said 95 percent of its energy needs come from its fixed and mobile networks, with "much" of the energy usage going towards PSTN/TDM technology. The company has, however, reduced its global energy consumption by 17 percent since 2017 – totaling some 420GWh – driven in part by the shutdown of its 3G network.

Openreach told *DCD* that fiber optic networks consume up to 70 percent less energy than copper networks. And, because they generate minimal heat, fiber reduces the need for cooling systems. The newer network also offers the benefit of longer lifespans and lower maintenance.

BT has partnered with IT asset disposal firms N2S and TXO for what its calling the Exchange Clearance Operation to remove, reuse, and recycle materials from exchanges. In 2023 alone, N2S said it aimed to support BT in extracting 200 tonnes of copper cable via recycling and the resale of redundant network equipment, as well as recycle over 2,000 tonnes of lead batteries. That figure reached more than 600 tonnes of copper in 2024. Part of that is being recovered via bioleaching, which uses bacteria to recover precious metals. Openreach told *DCD* that most of its waste will

be recycled rather than sent to other countries for spare parts in other legacy copper networks.

Exiting exchanges

The BT Tower exchange Openreach showed *DCD* around hosts some 6,000 PSTN and MPF lines. The tower is also home to some 28,000 lines of fiber with some 5,000 active circuits. While some fiber might be private lines for enterprise customers or for interconnecting exchanges, individual fiber lines can serve up to 32 residential customers on fiber to the premise.

The facility should be fully empty of customers by midnight on March 31, 2029; London's Marylebone Exchange, a mile west, is the receiving site for customers at the tower.

At the time of *DCD*'s visit to the tower, Openreach said it had contractors doing physical cable audits tracking the cable from end to end for around a month at that point. The processes just for this tower is expected to take several months at least.

"They are understanding what it is, and then they are reporting that back into our planning teams to understand where it goes," says Rory Lockyer, exchange exit, delivery professional for Openreach. "From there, we can then get a grasp of who is expected to be on that cable, what customers are we expecting to serve, and

where does that cable go?"

When asked if seeing the thousands of lines of copper on the MDF brings any kind of anxiety about getting everything and everyone out of the exchange on time, Lockyer says no, because many of those lines will have already been disconnected as people have already moved to fiber.

"As folks self-select onto fiber, you end up with wires here that you don't necessarily care about, so it becomes a lot easier to decommission," adds Chadha.

Openreach has already done preliminary work at many of the 105 sites that will be closed down following the initial pilot, including preventing new copper products being sold by CPs. Nesta Colburn, exchange exit manager at Openreach, said the company is starting the exit process by tackling "some of the most complex" sites.

While some parts of facilities, such as the MDF, fit a fairly standard design, exchanges vary dramatically in terms of size and layout, making each one unique in terms of the work needed to close it down. The company declined to tell *DCD* how much money it will be putting into the exit program due to its "commercial sensitivity."

The OHPs that Openreach is keeping reflect the geography of the company's fiber rollout. Those remaining will be large enough to fit all the Openreach and CP equipment required to serve customers



"Altnets have little confidence that Openreach will meet its initial target for exchange closures"

>>INCA

across the UK once the rest have been closed, the company says.

An OHP site hosting fiber services can be found inside the BT Tower, and is one of those facing the chop. The OHP areas of the tower exchange, akin to large meet-me rooms, feel similar to traditional white space, with one networking box hosting around 2,000 fiber customers, compared to the 100 lines a single similar-sized copper block could host on a copper network on the MDF downstairs.

In the BT Tower and other OHPs that will be closing, fiber customers will need to be migrated to facilities that are remaining open. Though the migration of this network will need far less manual intervention than the copper customers, some physical re-splicing of the fiber will be required in parts of the network.

All CPs present in closing exchanges will be offered equivalent space and power in other OHPs, Openreach tells DCD. Not all exchanges are head-end buildings, where fiber networks terminate, which the company said has been factored into some of the decisions around what sites Openreach will close in the program.

At BT Tower, Openreach is retaining the cable chamber vault where copper and fiber cables enter the building, and is constructing a small area to host passive network equipment.

"Because of some of the sensitive equipment we have in here, and some

of the contracting that we agreed when we sold these buildings, we are able to retain certain spaces," says Lockyer. "In [the BT Tower], we are planning to move all of our fiber equipment down to the cable chamber. Over the next 12 months, we're creating a space that will be used to passively manage all of our fibers, so there won't be any active equipment like head-ends in here.

"It will be a secure location that fibers can transit through to go on to their receive site."

Complications and regulations

While BT, via Openreach, still has a monopoly on exchanges, the shutdown will have a major impact on hundreds of communications providers across the UK.

Earlier this year, Openreach said its full fiber network is available to around 18 million homes and businesses, with some 6.5 million premises currently taking a service. Ofcom data suggests there are still around 5.2 million PSTN customers in the UK, totaling approximately 27 percent of all landlines. Openreach has said it is reaching an average of 85,000 new premises every week with its full fiber offerings and is targeting a total of 30 million premises by 2030. The company says its full fiber build-out will cost in the region of £15 billion (\$20bn).

But, like any data center migration, each exchange exit is a multi-step

process that involves multiple parties. The shutdown is a complicated process that is fraught with business, logistical, and regulatory hurdles at every step. While Openreach is a B2B firm serving communications companies, the impact on vulnerable residential customers will likely govern the final pace of the shutdown.

To begin the exit process, Openreach issues a stop sell notice to make sure it and any wholesale customers aren't connecting any new lines. Stop sells are triggered when a majority (75 percent) of premises connected to a particular exchange can be served via fiber.

Once a stop sell has been called, the companies can begin the slow process of migrating and disconnecting the thousands of cables in each exchange – a complicated operation that will involve Openreach working with its customers, and those communication providers working with each customer individually. Once everyone is fully migrated and disconnected – a process which will take years – equipment can be shut down and removed bit-by-bit.

BT and Openreach have said the exit of the simplest exchanges could be done within four years, but the process of exiting the most complex exchanges would begin up to seven years before their final exit date.

Chadha says: "We understand the pain of the journey, we're not blind to



it. Some of it is just really hard to move because you can't tolerate disruption. How do you make sure that you migrate the system that operates traffic lights, for example, without them stopping?"

He says some CPs have been resistant to change, and have to be convinced a migration is possible: "It has always started with 'this is why you can't do it,' but we've crossed the rubicon into 'how are we actually going to do that?'" Chadha says.

Lockyer adds: "There are going to be 1,000 different reasons why someone can't migrate. And we will come back with 1,001 answers or solutions on how to get them to migrate."

To help sweeten the deal, Openreach has a commercial offer that involves an incentive payment to CPs if they exit all exchanges on time. If CPs do not migrate all of their customers on time, they forego this payment.

It can be difficult to get a repeatable pattern on such migrations. In the traffic light example, the company has already come up against at least two separate

"While the focus is on the 108 that we closing, we need to be focusing on is the 89 that we're keeping. We worked out along the way how big of an expansion program this is."

*>> Nesta Colburn
Openreach*

traffic light networks, with different ownership and operating structures, operating with different network set-ups via different CP providers.

Colburn says: "It's not a cookie-cutter approach where we can say we've learned something here and we're going to apply that same process elsewhere; we start to find that we're doing bespoke solutions."

Despite this, it is hoped the answers to many core issues can be found during the pilot stage. An example Lockyer gives is the cable chamber being retained in the BT Tower; going forward, it is hoped any other retained vaults in buildings will be a quicker and easier process once the first one is complete.

The first stop sells were issued at two sites, in Salisbury and Mildenhall, in 2020 and 2021. Today, more than 940 exchanges – covering around eight million premises and some 44 percent of the company's fiber footprint – have received the notices, with another 226 exchanges set to be added the list next year. Customers on existing contracts can continue to use copper-based services, but will unlikely be able to renew their deals.

The rub is that an exchange can't be closed until the last customer has been disconnected. And that can't happen until a fiber alternative is made available and any vulnerable customers relying on a copper network have been switched onto a similar service. The fact that Openreach only deals with CPs and not the end

customer can make that process more difficult.

"Once you've got fiber, customers usually self-select onto that because it's a better product," says Chadha. "That takes the majority out of it, then you're left with the rest. And the economics, keeping this [exchange] open for that small proportion, ultimately, just doesn't work, for Openreach, for BT group, the UK, or industry."

He adds: "We don't have control of providing to the end customer, and that becomes really difficult if you're trying to close an exchange or the PSTN. If the CP hasn't done enough, how do you manage that end customer?"

He notes the challenge for CPs around migrating large enterprises. National and international chains are more likely to want a nationwide approach to switching to fiber, rather than the Openreach calendar where different exchange sites are exited at different times.

"Think about a hospital, or a supermarket, or a coffee chain. How do you start moving them all? If you're their CP, it's not an easy conversation," says Chadha. "The challenges of this stuff often come down to really hard commercial problems."

He tells DCD that regulation reform is needed to help ensure its mandated like-for-like services are available to every customer, especially in the residential space.

The company has historically struggled to gain access to multi-dwelling units, such as apartment buildings, in order to install new fiber lines. Landlords can be difficult to contact - or convince - to gain permission to install new fiber. Even if Openreach has the right to enter units to fix the copper network, it has no automatic upgrade rights to install fiber. Chadha notes Openreach also would like renters to have the right to install fiber without needing permission from the landlord.

He also questions why it needs to be Openreach providing the fixed-line fiber replacement for copper, when other fiber providers, or wireless networks, could offer a similar quality of service in many areas. Other countries have seen the copper-operating incumbent given the power to force-migrate customers to a

new service, which Openreach doesn't have.

"We all agree that every customer needs a decent connection. But as long as folks have a good connection, I think we need to be a little less obsessed about how it's provided and by whom," he says.

Connect Europe's report suggests operators should "receive strong support from national policy makers and regulators to assist them in their switch-off efforts."

The UK's telco industry watches and waits

While this is a copper-centric program, the exchange exit has a large impact on fiber providers, too.

As well as copper offerings, many exchanges offer Ethernet, Dark Fibre (DFX), and PIA (access to existing cable

ducts and poles) solutions. OHPs that are closing will see fiber services relocated to enduring exchanges, meaning equipment needs to be moved.

Openreach, created because the UK's wave of new telecoms didn't like how BT was abusing its 100-year monopoly and incumbent status in a modern deregulated market, still has an occasionally uneasy relationship with the rest of the UK's communication providers. The exchange exit is again causing the industry more worries.

The scale of the exit at each exchange will vary depending on the CP. Neos Networks, a UK fiber wholesaler, has previously said it has its optical equipment in around 550 Openreach exchanges to regenerate signals every 80 kilometers.

Even companies that aren't using Openreach's copper networks, such as some of the UK's new upstart fiber providers, or altnets, will have fiber

BAYNARD HOUSE AND FARADAY BUILDING

Along the banks of the river Thames and close to Blackfriars Bridge and station, it's hard to miss Baynard House. Once the site of Baynard Castle – destroyed in the great fire of London, the site was long home to warehouses. BT's four-story brutalist block was completed in 1979. As well as being an office and telephone exchange, it was also the home of the BT Museum for 15 years. Around 2012 it was the site of the largest rooftop solar deployment in the City of London and the second largest of a corporate building in the UK.

Notably for art lovers, it is home to The Seven Ages of Man by Richard Kindersley; an eerie cast aluminum sculpture commissioned by the Post Office and erected in 1980, featuring a totem pole of heads representing the different stages of life. Notably for cinema lovers, it is the roof of Baynard House that was used when actor Tom Cruise broke his ankle shooting a stunt for *Mission Impossible: Fallout* – a shot which made it into the final cut of the film.

Across the street is the much older

and more classically aesthetic Faraday Building. The five-story Neo-Georgian facility stands in stark contrast to its sibling across the street. The facade features nods to Hermes/Mercury, the Greek/Roman deity known as the messenger of the gods, as well carvings of telephones, switching equipment, and cables.

After previously hosting a building that was home to a society of lawyers, the land was first acquired by the General Post Office in the 1870s and saw a new building developed to house the Post Office Savings Bank. The building was converted into a telephone exchange around 1902; launching with some 200 subscribers, it had grown to some 10,000 within three years. The site was rebuilt for a dedicated automatic international exchange around 1933, with the new Faraday Building named after English electromagnetism scientist, Michael Faraday.

Notably, the construction of the Faraday Building obscured the riverside view of upper parts of St Paul's Cathedral, leading directly to regulations protecting views to and from the dome. ■

DRINKING IN THE SHADOW OF THE BT TOWER

Pubs with fun names are a staple of British culture. According to posts on the BT Telephone Exchange Enthusiasts Club Facebook group, it wasn't uncommon for the larger telephone exchanges to have bars in the building!

It is rare, though, that pubs are named in honour of digital infrastructure.

The Telephone Exchange in London Bridge opened in 2019. The building was erected in 1915 and was used as a wartime telephone exchange from 1940, which inspired the new name and décor.

Previously a Jamie Oliver restaurant, it was refurbished by brewery Fuller's, and "old telephones adorn the art deco walls, teamed with quirky furnishings and vintage leather seats," according to the company. Sadly, the site closed around 2023.

One of the few other pubs with a close history to our digital world that *DCD* is aware of, however, is unlikely to ever serve a pint again. In the shadow of BT Tower, at 2 Clipstone Street in Fitzrovia, lies the now-empty Tower Tavern.

The original pub on this site was first licensed in 1776 as the Bastard Arms. The origin of this name has been lost to time, but sources online say it was either named after the landlord, John Bastard, or in

honor of Henry Fitzroy, the bastard son of Charles II. It was renamed to Fitzroy Arms in 1826.

That original building was demolished and the new drinking establishment became the Tower Tavern around 1970, a few years after the opening of the BT Tower (then the Post Office Tower).

A flat-roofed cube of white tiles, it isn't the prettiest venue. A local community Facebook group once labelled it "the worst looking pub in London."

It closed around 2021 in the wake of Covid-19. The University of Westminster owns the property and reportedly chose not to renew the lease after it expired.

The existing building was set to become a new student hub as part of a wider redevelopment of Cavendish Block into a new student experience center, but its current status is unclear. A report from late 2024 said the University was "exploring options" for the former pub.

DCD heard the pub has been bought by the same buyers of the BT Tower, but the university told this publication is it "not selling the Tower Tavern pub."

We have followed up about its plans for the site, and what might become of the delightful sign above the door. ■

that idea has since been withdrawn. A 2025 survey by Neos suggested rerouting networks due to exchange closures will cost companies an average of £1.4m (\$1.88m) each – a hefty sum when many are already operating at a loss in an attempt to scale.

And while the enduring OHPs stand ready to host the increased amount of fiber going to each site whether they are ready to host more IT equipment is less of a sure thing.

CPs and altnets house racks of network gear at exchanges in the MUAs, and Openreach is offering like-for-like space and power in its OHP exchanges. It also pays compensation for CPs relocating racks from one exchange to another. That offer, though, could see those continuing OHPs filling up fast, and could require a power upgrade.

Racks in the tower exchange will be moving to Marylebone, along with the networking gear of two other sites. Elsewhere in London, exchanges in Wandsworth and Streatham will be consolidating into an enduring facility in Balham, which Openreach says will likely need a power upgrade to accommodate all the new gear.

"While the focus on this program is on the 108 exchanges we're closing, we need to be focusing on the 89 that we're keeping, because those are the ones that are going to be here forever," Openreach's Colburn says.

The company is now working with CPs to establish if moving all the equipment is really necessary: "If a customer has a point of presence in the tower exchange, North Paddington, Paddington, and Marylebone, do they actually need four tiles in Marylebone or could you make do with two, because all of the current equipment is only operating at about 50 percent?," Colburn says.

Cost savings from using less space and, potentially, less power in consolidated sites might be enough of a carrot to encourage some CPs to move, it is hoped.

The wider BT Group is already investing in some upgrades. In its 2025 annual report from May, BT said it had installed and upgraded cooling plants in its local exchanges, saying it had invested more than £9m (\$12m) on adiabatic

equipment in exchanges and use other Openreach infrastructure impacted by the exit program.

Though moving racks of one MUA into another at an enduring exchange isn't as complicated as the copper migration and closure, the Independent Networks Co-operative Association (INCA), which represents UK ISPs, has warned that the exchange exit could cause problems for the altnet fiber providers, both directly and indirectly.

Many companies rent duct and pole space from Openreach (a service known as Physical Infrastructure Access, or PIA) to save building out their own. INCA has warned that although some

PIA infrastructure might no longer be required by Openreach once an exchange closes, it could still be important to altnets.

INCA has also bemoaned the fact that "effectively, it is up to Openreach's customers to shoulder the burden of rearranging their networks at their own cost – except where Openreach has specific contractual obligations to cover costs for its customers, which is extremely limited."

Openreach has said customers will need to rearrange their network to hand over at an alternative location if an exchange is being exited. The company briefly proposed allowing PIA users to acquire this unwanted infrastructure, but

cooling systems upgrades in the last year.

And even if companies don't have hardware in exchanges or are users of PIA services, the facilities act as 'meet-me' points and are key for the third-party supply of backhaul and other services to altnets. Openreach dark fiber services generally connect back to exchanges, and some altnets or their suppliers use this DFX service for backhaul - which could be shut down in some areas if an exchange is closed.

Openreach told *DCD* that the network in the exchange area is not impacted by the exchange closing, so the network element for PIA CP that is using Openreach duct or poles in the area will not be impacted. However, the company noted that if an alt-net or PIA CP has opted to route their network into a closing exchange, it is "their responsibility to re-route/rearrange this to another exchange or another location of their choosing before the exchange closes."

The industry group has noted that altnets are also concerned that increased demand at Openreach's enduring exchanges could result in a shortage of space and power when relocating. There could also be challenges in ensuring there is enough inter-exchange capacity if ducts between two remaining sites are full. Current PIA rules don't compel Openreach to build more duct space or remove idle copper, potentially preventing new altnet fiber taking its place. *DCD* understands BT is working to ensure there is sufficient space and energy at all exchanges.

"The closure of approximately 80 percent of Openreach's exchanges will significantly impact the current and future design of altnet networks and will result in significant costs for altnets to change their networks to mitigate the closures," INCA said in a report. "If Openreach is only required to provide access to space and power in an enduring exchange on the principle of 'where available,' this could create significant capacity constraints."

Not all companies will be impacted equally. Some altnets, like CityFibre, have been busy building out small exchanges over the UK. The company has at least 100 'FEX' deployments, with most being small containerized pods with a handful of racks.



Credit: Sebastian Moss

The closing of exchanges will also likely have an impact on local wireless services. Many exchanges host cell towers and other wireless equipment for BT/EE and others; the closing and potential redevelopment of those sites could mean the wireless infrastructure will have to be removed. Cellnex's On Tower took BT to court in 2024 over the latter's attempt to terminate a roof lease early at an exchange due to be decommissioned; the case is still ongoing at time of writing. Both companies declined to comment on the matter.



The view from the tower

Get out or pay up

Though BT says there are operational and sustainability reasons for the copper switch-off, it also faces a very real financial penalty if it misses the looming 2030 deadline to exit exchanges. Fittingly for London, a city with an ongoing rental crisis, BT and Openreach are very wary of their landlord putting up the rent.

BT sold the majority of its real estate sites to real estate firm Telereal Trillium in a sale-leaseback deal back in 2001 for £2.38bn (around £4.45bn/\$5.27bn today, adjusting for inflation). The company was aiming to reduce its debt pile, which was hovering at around £30 billion (\$40.19bn) on the back of the build-out of its then cutting-edge 3G network. Today, amid an ongoing 5G build-out, BT's debt is just under £20 billion (\$26.8bn).

Openreach tells *DCD* the six million sq ft (557,420 sq ft) portfolio was leased back to the BT group on "very favourable" terms, with minimal rental increases in the intervening years. However, once the current leases end in 2031, the telco faces a stark choice: Be fully exited from a facility, lease the property again for a minimum of ten years (until 2041) at today's market value, or buy the property back, again at 2025 valuations. This applies to any of those thousands of exchanges from which BT is not fully exited.

"2031 felt like a long way away [in 2001]. Then, all of a sudden, when we got to 2020, we were starting to think about this lease break coming up in 10 years," says Colburn. "And now we've really started to ramp up."

Property value was a major factor in identifying the first 100 buildings to close before that 2031 deadline. Given that many exchanges are located within inner city metros - including central London - BT Group can ill-afford to risk being locked into thousands of expensive long-term contracts at sites it doesn't want. This is "definitely a big deal" for the company, its execs told *DCD*.

Many of the facilities in the initial tranche of 108 exchanges being closed are located in London; these include Baynard House near Blackfriars, as well as facilities in Bayswater, Wapping, Wandsworth, Streatham, Kensington, Southwark, Shoreditch, Pimlico,



Paddington, Monument, New Cross, and Mayfair. Further afield, exchanges are due to be closed in the centers of towns and cities, including Sunderland, Liverpool, Birmingham, Glasgow, and Edinburgh, where property prices are higher than in rural areas.

Aside from the BT Tower, only one other exchange in the initial 108 exchanges isn't currently owned by Telereal; the Skyport exchange near Heathrow Airport, owned by UK aviation regulator, the CAA.

When asked if all the enduring OHPs will be owned by BT Group or Telereal Trillium, Openreach said OHP locations will be housed in "enduring sites where Openreach can agree terms with the landlord."

DCD reached out to Telereal Trillium to discuss its plans for the Openreach/BT exchange portfolio once the telco exits the sites, but the company didn't respond. An Openreach spokesperson said the landlord would likely sell the buildings off to developers – something it has done with other exited BT properties. BT Group declined to comment on its real estate plans.

"We're freeing up great properties like the tower to create additional economic value from becoming an amazing hotel that people are going to travel for," says Chadha.


"Right across the country, what can you unlock by freeing up those exchanges that are often near high streets?"

Could they make good data centers still? Potentially. Some companies in other countries have repurposed legacy

exchanges into colocation facilities. They can't offer the same kinds of densities as modern facilities, and operators have to contend with the fact power is usually distributed in 48 VDC power rather than AC, but they do offer ready-made capacity in key locations.

But the fact that many such properties are in central locations in major metros means they are prime redevelopment opportunities into apartments – most exchanges exited in the last 20 years have been turned into residential accommodation.

Though it is part of the wider exit plan, the BT Tower was not actually part of the Telereal Trillium deal. The BT-owned site was recently sold to a hotel chain for redevelopment. Both Openreach and BT are involved in their own respective exit programs – the tower is also a broadcast operations center and has a separate BT-operated data center in the complex – but on different terms to the rest of the portfolio.



The future legacy

Referring to all the PTSN gear at the BT Tower, Chadha says: "In 2027, all of this has to disappear, and that date has to stick."

On the wider exchange shutdown, Chadha reports mixed success from the three pilot sites: "My confidence today is that we've got two of those three that look very likely that they are going to close and that we actually will succeed," he says. "The third one's difficult, but we'll see where we get to."

It might be a long time off, but what happens after BT and Openreach close all those exchanges? The group will have a smaller, more energy-efficient network and real estate portfolio for sure. But does that open up new opportunities for a new Edge play?

David McKean, BT's specialist tower division director, recently told *TelcoTitans* that the company has been considering whether it can repurpose some of its enduring exchanges and its tower portfolio for data center use. The telco also recently posted a job opening for an Edge computing architect, suggesting moves could be afoot.

BT Group owns at least 200 radio masts and towers in Britain – more than a dozen of which are large concrete towers in the style of London's BT Tower.

The group's total cell tower footprint, used by its EE unit, is unclear; the company said EE's 3G switch-off involved retiring technology at more than 18,000 mobile sites, with more being rolled out to support the group's 4G and 5G network.

BT previously sold around 220 of its

towers to Cellnex in 2019 for £100 million (\$134m), with the telco remaining a customer at a number of the acquired sites.

Several European operators, including CRA in Czechia, Digita in Finland, Telecentras in Lithuania, LVRTC in Latvia, and Cellnex in the Netherlands, as well as HKBN in Hong Kong, have deployed data centers at TV and radio tower sites. Edge data centers at cell tower sites are more common, with the likes of Cellnex, American Tower, and SBA having deployed data centers at cell sites.

The relationship between BT and Openreach could well change too. Allison Kirkby, BT's CEO, has previously hinted that Openreach could be spun off into a completely separate company once the group has completed its fiber roll out.

But will Openreach and BT meet their punchy targets to exit the 100 or so initial exchanges, let alone the thousands it is targeting long term? Despite delays to closing the initial three (already downsized from the original five sites), Openreach seems to think so. Not all are so convinced.

"Altnets have little confidence that Openreach will meet its initial target for exchange closures and significant delays may be incurred in the overall programme," INCA said in its report, "similar to the delays being faced by Openreach on its PSTN switch-off program."

We are reaching the end of the analog era. This writer just hopes whoever takes over Baynard House lets us keep playing football there once the last copper wire is pulled out. It really is hard to beat that view, even if our win record in the shadow of the tower isn't quite as spectacular. 

COPPER GETS THE COLD SHOULDER

For more than 100 years, copper has formed the backbone for much of the world's connectivity. And even after fiber started being used for communication networks in the 1970s, the orange metal would continue to rule for decades.

But today, copper network shutdowns are happening around the world amid ubiquitous fiber connectivity.

The European Commission has set a 2030 copper shutdown date for its member states as part of its push towards gigabit networks. Few have fully managed the feat to-date, though Connect Europe predicts 21 PSTN networks in Europe will have been decommissioned by 2030.

BT and Openreach won't be the first in Europe to call time on copper; Norway's and Spain's incumbent telcos have already completed the majority of their switch-offs.

Telefónica recently completed its switch-off in Spain, reducing the number of exchanges from more than 8,500 nationally to around 3,000 fiber-centric sites. Telefónica first shut off copper networks in Sant Cugat del Vallès (Barcelona) and Torreldones (Madrid) back in 2015; the first exchanges that hosted other providers were shut in Clot (Barcelona) and Hermosilla (Madrid) in 2021. The last tranche of 661 copper exchanges were closed in May 2025.

The UK won't be the last though; Germany is yet to make firm plans to get off copper – though Deutsche Telekom is

known to be interested – while Greece is yet to announce a switch-off date.

Telecom Argentina started its copper shutdown in 2023; Ethiopia's Ethio Telecom, Norway's Telenor, Ireland's Eir, Belgium's Proximus, France's Orange, Portugal's MEO, Sweden's Telia, and Monaco Telecom are others that have or are in the midst of copper retirement. Proximus is said to be mulling the sale of some 500 properties as part of its network switch-off.

"We believe that fiber networks are fundamental to the digital transformation of Europe," FTTH Council Europe President, Roshene McCool, said in January 2025.

"Phasing out copper networks for fiber infrastructure will lower energy consumption and reduce overall operating costs, therefore making a great contribution to the achievement of the EU's Digital Decade objectives"

In the US, the telephone exchanges are owned by a number of incumbent local exchange carriers (ILECs); the local telephone companies that held the regional monopoly on landline service before the market was opened to competition in 1996.

Companies with ILEC networks include Verizon, AT&T, Lumen, Brightspeed, and Zipy. A 2015 report from FierceWireless suggests there could be as many as 30,000 Central Offices across North America.

While most telcos have long ago sold off their core data center assets, telephone exchanges offer a way for incumbents to potentially get back into the Edge sector if they hang onto and repurpose assets, rather than sell them off. Telecoms consultancy firm STL Partners predicts there will be as many as 1,800 "network Edge" data centers globally by 2028 – up from just over 800 today.

Zipy has upgraded its copper network to fiber and used the excess real estate to start offering retail colocation services. AT&T, meanwhile has signed several sale-leaseback deals for a number of its copper exchanges ahead of its own network retirement plan.

Frontier – which is being acquired by Verizon – previously announced a deal to let AT&T deploy equipment for its 5G network in its former copper COs. Lumen has also reportedly pivoted some of its old Central Offices to data center colocation. Frontier also offers what it calls 'Edge Colocation' at more than 2,500 locations, including its Central Offices. Brightspeed – acquired from legacy Lumen assets – also offers colocation services from its facilities.

January 2025 saw AT&T seal a \$850 million sale-leaseback with real estate firm Reign Capital for 74 underutilized Central Offices, following a similar deal with Real in 2021 for 13 other properties. The company, which plans to retire its copper network completely by 2029, said the deal comprised a "small portion" of its copper portfolio. 



PROVEN AND SCALABLE DATA CENTRE SOLUTIONS

From digital energy management solutions to backup power and renewable integration ABB's proven and scalable data centre solutions help our partners boost energy efficiency, save operational costs and protect uptime.

**ENGINEERED
TO OUTRUN**






Sponsored by

Life Is On

Schneider
Electric

The Critical Power Supplement



Energizing tomorrow's data centers

INSIDE

Nuclear option:

>Colocating with
power plants

Growing capacity:

>Using AI to improve
the grid

Free flowing:

>Flow batteries and
data centers

A man and a woman are in a server room. The man, wearing glasses and a blue shirt, is holding a laptop and looking at the screen. The woman, wearing a grey blazer, is looking at the laptop screen. They are standing in front of server racks. The background is dark with some green light from the server racks.

Power quality solutions engineered to prevent data center downtime

Voltage sags and harmonics are the leading threats to uptime. Schneider Electric's **PowerLogic™ DVR** and **PowerLogic™ AccuSine PCS Plus** deliver precise power correction and ensure compliance to IEEE 1668 and IEEE 519 standards.

Boost resilience. Protect equipment. Stay operational.

[Download the Power Quality for Data Centers eGuide](#)



Sponsored by

Life Is On

Contents

34. Colocation Meets Regulation

The what and why behind FERC's rejection of AWS and Talen's behind-the-meter agreement

38. Quality issue

Why power quality is the unsung hero of the AI-driven data center

40. From Years to Months

Creating an AI fast lane for grid connection

44. Go with the Flow

Are flow batteries the answer to the energy storage needs of data centers?



Powering up

As the era of the gigawatt data center approaches, securing a reliable power grid connection has never been more critical.

Global electricity consumption from data centers is projected to more than double, according to a recent International Energy Agency (IEA) report, with AI cited - perhaps unsurprisingly - as the major driver.

The IEA report projects that data center consumption globally will grow to 945TWh per year by 2030, from 415TWh in 2024.

With this in mind, the fight for energy resources is likely to ramp up in the coming years, so operators are looking for novel solutions.

Colocating data centers at power stations has been touted as one potential solution, with behind the meter agreements potentially allowing firms to access power directly from the source before it even hits the grid.

However, one attempt to set up such an arrangement, AWS's plan to take power from Talen's nuclear plant in Pennsylvania, US, to supply a newly-acquired has run into regulatory difficulties. Zach Skidmore takes a closer look at the problems that occur when colocation meets regulation.

Elsewhere, AI could also be part of the solution, with several companies developing solutions

that deploy artificial intelligence to look for areas where the grid could get more efficient. In this supplement, we profile one of these companies, GridCARE, an US vendor using AI to map available grid assets that can be used to relieve constraints in real time.

By combining grid data, asset visibility, and scenario modeling, the company believes it can unlock significant extra capacity.

As more and more renewables come into play in power systems across the world, battery energy storage systems, better known as BESS, are becoming a staple of new data center designs. BESS enable energy to be stored so that it can be used at times when the sun isn't shining or the wind has dropped.

While the need for BESS is not disputed, the type of battery that should sit at the center of these systems is up for debate, with the dominant lithium-ion chemistry not necessarily the most suitable for the task, not least because a fire or explosion could have big implications when these batteries are located in multi-billion dollar developers.

A growing number of vendors believe liquid flow batteries could be a safer, more efficient alternative to lithium-ion for data centers, but, as we discover, significant hurdles must be overcome if they are to become a commercially viable option in the digital infrastructure space.

Colocation Meets Regulation



Zachary Skidmore
Senior Reporter
- Energy and
Sustainability

*The what and why
behind FERC's
rejection of AWS and
Talen's behind-the-
meter agreement*

In March 2024, AWS made headlines with the acquisition of the 960MW Cumulus data center in Luzerne County, Pennsylvania, from Talen Energy. The deal was not just another major data center acquisition; unbeknownst to the sector, it was set to have wide-ranging implications on how data centers are powered.

As part of the acquisition, the companies signed a behind-the-meter agreement that would have delivered 480MW of power directly to the data center from the 2.5GW Susquehanna Nuclear Power plant, located adjacent to the facility.

What followed, however, was a regulatory standoff that would make waves across the energy and data center markets, sparking unanswered questions about how regulators and utilities can best manage the explosive growth of power-hungry AI and cloud infrastructure.

Behind-the-meter

Unlike a grid-based or front-of-the-meter agreement, a behind-the-meter agreement involves power being supplied directly to an installation via a direct connection, known as a “private wire.”

As a result, behind-the-meter agreements bypass the main electricity meter for some or all of the consumer’s power needs. For data centers, this can be especially beneficial, resulting in reduced energy costs, a faster connection time, and a more reliable source of energy.

While there have been many instances of behind-the-meter agreements in the data center sector, the AWS-Talen agreement differed in both scale and choice of energy. Unlike previous instances, often utilizing onsite renewables, the AWS deal involved a regional key generation asset, which provides consistent and reliable power to the grid.

As a result, to secure the go-ahead, PJM Interconnection, the regional transmission operator in charge of the utility services in the state, had to apply for an amendment to the plant’s existing Interconnection Service Agreement (ISA), permitting the increased power supply.

However, rather than the swift approval the companies hoped for, two major utilities that operate in the region, Exelon and American Electric Power

(AEP), vehemently opposed the amended ISA, submitting a formal objection to its provisions.

Exelon chief operating officer Michael Innocenzo revealed the rationale behind the opposition, stating: “Our main concern [was] that colocation shouldn’t be a way to bypass distribution or transmission charges. These facilities still rely on the grid, and we’ll need to invest in upgrades to support them, whether they’re inside or outside the fence.”

Due to the protest, the ISA proposal was escalated to the Federal Energy Regulatory Commission (FERC), which in November 2024 rejected the amended ISA in a two-to-one vote, stating that the parties did not make a strong enough case

to prove why a special contract should be allowed in this instance.

“The burden is on PJM to show that nonconforming provisions are necessary, and PJM failed to carry that burden,” read a FERC statement at the time.

FERC reaffirmed its decision in April 2025 after a rehearing request from Talen, ultimately scuppering the agreement, leading to the companies having to explore further options. The implications of the FERC ruling are notable in setting a high legal bar for future behind-the-meter deals. But, despite the ruling, debate continues over why such barriers exist, highlighted best by then-chair of FERC Willie Phillips, who, amongst his peers, was the only dissenting voice, voting in favor of the amended ISA.

A missed opportunity?

The denial by FERC of the AWS-Talen deal was rooted in the belief that the parties involved did not make a strong enough case to prove why a special contract allowing for expanded “behind-the-meter” power sales should be permitted.

For Phillips, however, the case was never about approving the “perfect scenario” for colocation, but rather an opportunity for a test case of how colocation could function in the future.

In his support of the ISA, Phillips contended that it

“One of the worst things regulators can do is become a bottleneck for innovation. That’s what concerned me most. The PJM case could have been a learning moment. Instead, we missed a chance to get ahead and evaluate colocation in a real-world context,”

>> Willie Phillips

would have provided the ability for the generators, consumers, and broader regulatory community to better understand the potential of colocation, acting as a “pilot” for more flexible load models, which he considers a necessity to meet the staggering demand of data center growth.

“I saw the PJM case as a one-off opportunity to pilot colocation. We could have required updates on reliability risks, encouraged targeted investment, and learned from how it performed,” Phillips says.

In his statement explaining his dissent, Phillips went on to contend that, rather than just providing free rein, as feared by objectors, he would have ensured that PJM submitted regular informational filings to provide transparency into the arrangement’s operations over time.

In addition, he said that the deal would have allowed PJM to go through a further stakeholder process for tariff revisions and decide on generic next steps. Therefore, by rejecting the ISA, Phillips argued that FERC essentially “rejected protections that the interconnected transmission owner says will enhance reliability while also creating unnecessary roadblocks to an industry that is necessary for US national security.”

The centrality of data centers as a “national security asset” clearly played in the mind of Phillips in his decision-making. Data centers, after all, are some of the most sought-after infrastructure for national governments, due to their ability to attract capital and expertise to a nation’s coffers.

In turn, Phillips believed that regulators should seek to support the sector through novel deals, as the skyrocketing demand offered an opportunity to act as “a catalyst to modernize and upgrade our transmission system” and support the projected demand growth across the country.

Consequently, Phillips’ decision centered greatly around the notion of whether regulators should act as a roadblock or a facilitator of change. “One of the worst things regulators can do is become a bottleneck for innovation,” he says. “That’s what concerned me most. The PJM case could have been a learning moment. Instead, we missed a chance to get ahead and evaluate colocation in a real-world context.”

The learning moment could also have

supported greater alignment between the regulators, which all play a crucial role in the process. Made even more important due to the scale and speed of demand growth emanating from the sector, at a rate “not seen since the Industrial Revolution,” Phillips continued.

Therefore, for Phillips, the lack of regulatory clarity within the industry as it pertains to how to power these massive facilities is proving a thorn in the side of progress, which in turn hurts planning procedures that could provide a streamlined process for large load connections while avoiding the worst impacts, namely shifting costs onto consumers.

Pay for play

The issue that ultimately scuppered the AWS-Talen deal was the notion that the companies, via a private wire, were shirking the responsibility to pay their fair share of infrastructure costs to upgrade the network.

As Michael Innocenzo puts it, it was never a problem with the idea of colocation itself. “To set the record straight, we’re not against colocation in any way,” he says. “If a large data center can colocate with a generator and get online faster, we fully support that.”

Instead, Exelon and AEP sought to ensure that “colocated facilities still use grid services and should not be exempt from paying for them. They are network

“We need regional coordination, and I would support a national plan to integrate generation and transmission for data centers. They’re not like other loads—they’re something entirely new, and we have to plan accordingly.”

>> Willie Phillips



load and should be treated accordingly,” Innocenzo argues.

The companies refuted the claim made in the proposal that colocation would remove all grid impact. Arguing that even if the load directly connects to the generator, the generator is still reliant on the grid itself for stability and voltage regulation. Under current US laws, the generator is not expected to pay for these grid services, with the load user (ratepayer) footing the bill. However, under the proposed ISA, AWS, as the ratepayer, would avoid the responsibility to foot the bill.

Therefore, for Exelon and AEP, it was all a “matter of affordability,” says Innocenzo, as while the companies want to support data center growth, they also want to avoid “shifting costs from high-usage customers onto residential ratepayers.”

This is an issue that Phillips agreed upon, viewing the importance of affordability as “critical.” However, he argued that this falls on the regulators to ensure proper planning and create provisions to prevent cost shifting onto residential and industrial customers who already fund the grid’s backbone.

What is even more crucial is wholesale investment in transmission and distribution infrastructure to meet the expected demand. Here is where Phillips has some concerns: “You don’t want a situation where investment in AI infrastructure outpaces the investments we need in transmission and generation. That’s where we risk falling behind.”

As a result, we are increasingly seeing utilities and state legislatures bring forward new rules on large-load customers and their contribution to

"To set the record straight, we're not against colocation in any way. If a large data center can colocate with a generator and get online faster, we fully support that,"

>> Michael Innocenzo

funding grid infrastructure projects. Notable recent examples include Oregon, which recently passed a bill that would provide regulators the ability to ensure that data centers and other large loads cover a fair share of the costs associated with new power plants and transmission lines used to power their installations.

New Jersey and Ohio have also proposed similar bills to ensure this, marking a prevailing trend of regulators seeking increased protections for ratepayers over who bears the costs of generation and distribution.

Is this the end for behind-the-meter?

Since the rejection by FERC, Talen and AWS have reimagined the agreement, with it moving from behind to an in-front-of-the-meter arrangement. The 17-year PPA will see Talen supply AWS with 1.92GW of power, ramped up over the next seven years, with the power provided through PJM. This reflects a broader move within the sector, with both Talen and nuclear energy generator Constellation

indicating their intention to focus on grid-based arrangements going forward.

Despite this, Phillips still believes that under the correct circumstances, colocation can be a powerful tool, especially for AI and hyperscale cloud deployments seeking to scale quickly. To ensure this is done effectively, however, Phillips contends that modernization is required across the broader regulatory landscape.

Phillips cited the 206 proceedings ordered by himself in February as a reflection of this urgency for reform, which he says could force through faster timelines for decisions on matters such as colocation. The proceeding aims to address regulatory gaps in how co-located large loads, particularly AI-driven data centers paired with on-site generation, are governed under the PJM Tariff.

In the current proceeding, FERC is evaluating whether PJM's current rules are unjust or unreasonable and whether new provisions are needed to ensure transparent, non-discriminatory treatment of co-located load configurations. The ultimate goal is

to provide regulatory certainty, support infrastructure investment, and ensure grid reliability amid a sharp rise in large-scale data center demand.

"What I liked about the 206 proceeding is that it boxes FERC in - it says, you have six months. You can't take two years. And I'm hopeful we'll get regulatory certainty from FERC very soon," says Phillips.


As a result, while colocation appears to be very much on the back burner, it is unlikely to disappear completely, especially as data center growth continues.

A rallying call

While the future of behind-the-meter agreements, especially for large hyperscale facilities, remains uncertain, Phillips contends that what's key now is developing a clear and consistent regulatory regime to ensure that the concerns voiced over colocation, namely, costs falling on ratepayers, are accounted for.

Consequently, for Phillips, a key factor in supporting colocation is greater coordination and planning across federal and state bodies. "This can't be business as usual," he argues. "We need regional coordination, and I would support a national plan to integrate generation and transmission for data centers. They're not like other loads—they're something entirely new, and we have to plan accordingly."

Subsequently, Phillips stated that "If I were still a regulator, I'd focus on three things: reliability, affordability, and boldness. We can't wait for a crisis before acting. We must lead—be proactive, not reactive—and lay the foundation for this next era of energy transformation led by data center demand." ●



Getty Images

Beyond uptime: Why power quality is the unsung hero of the AI-driven data center

SVP of digital power business at Schneider Electric, Arnaud Cantin, explores why it's not enough to source more power for the AI era – but power fit to fuel the data center machine

DCD: The data center industry is experiencing an unprecedented boom, largely driven by the rapid growth of AI. While much of the conversation focuses on capacity and cooling, you've been a vocal advocate for a topic that's often overlooked: power quality. Why is this so critical right now?

Arnaud Cantin: It's a crucial conversation. For years, the industry has rightly been obsessed with uptime – with Tier ratings and redundancy as the gold standards. We've mastered the art of keeping the lights on. But the nature of the load is changing dramatically.

Think of an AI data center not as a building, but as a Formula 1 (F1) engine. For years, we focused on ensuring it always had fuel (uptime). But now, the performance is so high that the quality of the fuel (power quality) is paramount.

You wouldn't put dirty, low-octane gas in an F1 car and expect to win. The relentless surge of AI, with its demand projected to grow at a staggering CAGR of 25-33 percent through 2028, puts an entirely new kind of stress on our electrical infrastructure. The greatest threat to resilience today may not be a complete blackout, but the constant, insidious degradation of power quality, and consequently noncompliance with grid codes.

We saw a case with a hyperscaler where a minor, sub-cycle voltage sag from the utility – far too short to trigger the UPS – was enough to cause data corruption across a whole rack of GPUs running a critical AI training model.

There was no outage in the traditional sense, but the cost to restart the multi-day

training job and the potential for flawed results was in the seven figures. That's the insidious nature of this problem. In this new era, focusing only on uptime is like building a fortress but ignoring the quality of the water or air. It's a fundamental vulnerability.

DCD: That's interesting. Most operators are focused on the grid coming in. Are you saying they're creating a significant part of the problem themselves? How can a brand-new, high-tech server be 'dirty' from an electrical standpoint?

Arnaud Cantin: Certainly. It's a two-sided challenge. On one side, you have the external grid. With the increasing integration of intermittent renewables and the rising frequency of extreme weather events, the power being delivered to a facility's doorstep is less stable than ever.

This is the contaminated fuel being delivered from the pump. Power quality issues are now the most common cause of major outages. On the other side, you have an internal threat that the data center creates itself. The very components that power AI – high-density server racks, advanced GPUs, and modern cooling systems with EC motors are non-linear.


These loads are sensitive to external voltage fluctuations (sags/swells) and at the same time generate significant electrical noise in the form of harmonic distortions, which contaminates the facility's network, leading to overheating, premature equipment failure. In our F1 analogy, this is like sludge building up inside the engine itself, created by its operation, reducing efficiency and causing wear.

DCD: This internal pollution doesn't just

stay within the data center walls, does it? This is where the conversation turns to compliance.

Arnaud Cantin: Exactly. Data centers are no longer just passive consumers of energy – they're significant players on the grid, accounting for up to three percent of global electricity consumption. This internal electrical noise travels back out, and utilities are taking firm notice.

Grid codes, like IEEE 519 in North America and EN 50160 in Europe, are no longer simply recommended, but are being enforced with increasing stringency by



In the AI era, mastering power quality isn't just about defense – it's a competitive advantage

*>> Arnaud Cantin,
Schneider Electric*

utilities across the globe. Operators who fail to keep their harmonic distortions in check are no longer just risking their equipment; they are facing significant financial penalties for destabilizing the local grid. Being a good grid citizen is now a matter of both operational necessity and financial prudence.

DCD: Let's be frank, this sounds complex and potentially expensive. For an operator who hasn't experienced a major power quality failure yet, how do you justify the upfront investment in advanced monitoring and active filters? Why isn't a top-tier UPS and a generator enough anymore?

Arnaud Cantin: That's the critical question. A UPS is essential, but it is not a panacea for all power quality issues. To build a truly resilient electrical architecture, we need a modern, multi-layered defense. The first step is to see the invisible. You cannot fix what you can't see.

Advanced, Class A power quality meters installed at key points provide the high-resolution visibility needed to diagnose these issues in real-time. These meters follow strict international standards that define precisely how measurements are taken – using consistent, validated algorithms – so you can trust the data when making critical decisions.

DCD: That brings us to the crucial first step. Before you can even begin to "see the invisible" or "cleanse the power," you need accurate data. How does robust power metering lay the essential groundwork for achieving better power quality in a data center?

Arnaud Cantin: Absolutely. Think of power metering as the vital diagnostic system for our F1 engine. Without it, you're driving blind. High-resolution, granular power metering at various points – from the utility incoming feed to main distribution panels, down to individual racks and even specific critical equipment – provides the foundational data.

It's not just about energy consumption for billing – it's about understanding the nuances of current, voltage, frequency, power factor, and crucially, harmonic distortion levels. By continuously monitoring these parameters, we can establish baselines, identify deviations from ideal sine waves, pinpoint the exact source of internal electrical noise, and detect subtle external grid disturbances that might otherwise go unnoticed by a

standard UPS.

This detailed, real-time data is indispensable for diagnosing power quality issues, trending performance over time, and validating the effectiveness of any mitigation strategies. In essence, comprehensive power metering transforms a reactive approach to power quality into a proactive, data-driven strategy, allowing operators to make informed decisions and truly master their electrical ecosystem. But seeing the problem is only half the battle. You must actively cleanse and correct the power.

DCD: This is where solutions like Schneider Electric's AccuSine range come into play. How do they specifically address that compliance challenge?

Arnaud Cantin: This is precisely where we focus our innovation. Once you've identified harmonic distortion, you need to eliminate it. Our PowerLogic AccuSine PCS Plus Power Correction System acts like noise-canceling headphones for the electrical system. To use our engine analogy, it's the sophisticated fuel filtration system that cleans out the sludge in real-time.

They dynamically monitor the network and inject an opposing current to cancel out the harmonic distortions created by the IT loads. This 'cleanses' the power, which not only protects sensitive downstream equipment and improves overall efficiency but, crucially, ensures compliance with utility grid codes.

The PowerLogic AccuSine PCS Plus Power Correction System is a direct response to this challenge, giving operators the tool they need to mitigate harmonics at the source and avoid those costly penalties from utilities. It's about moving from a passive to an active stance on power quality.

DCD: So the AccuSine range is about actively ensuring you're not polluting the grid, while other solutions protect you from the grid's instability?

Arnaud Cantin: A perfect summary. You use a solution like the PowerLogic AccuSine PCS Plus Power Correction System to ensure your facility is a responsible grid citizen. Then, to shield yourself from external grid volatility, you can deploy a Dynamic Voltage Restorer (DVR) for brownout conditions.

A DVR sits at the front end of a facility and acts as the ultimate shield, instantly

correcting for voltage sags and swells without ever engaging batteries. It ensures a perfect sine wave reaches your critical load, regardless of fluctuations from the grid.

DCD: It seems the final piece is bringing all this information together.

Arnaud Cantin: Precisely. The final layer is a sophisticated power management system. A platform like our EcoStruxure Power Operation aggregates data from all the monitoring devices and mitigation equipment into a single pane of glass.

This allows operators to shift from being reactive – chasing phantom failures after they occur – to being proactive.

They can use analytics to spot negative trends, validate that their AccuSine filters are ensuring grid code compliance, and continuously optimize the health and efficiency of their entire electrical infrastructure.

DCD: To conclude, what is your key message to data center operators navigating the AI era?

Arnaud Cantin: Ultimately, operators are at a crossroads. They can continue to react to a power environment that is fundamentally changing, treating these phantom failures as the cost of doing business. Or, they can become active masters of their electrical ecosystem.

In the AI era, mastering power quality isn't just about defense – it's a competitive advantage. The winners will be those who ensure the health of every single watt. The conversation has moved beyond uptime – it's now about the holistic health of your power.

To learn more about Schneider Electric's power quality solutions, visit the website. You can download Schneider Electric's latest e-guide about power quality for data centers here. ●

>> <https://www.se.com/ww/en/download/document/998-23590151/>



From Years to Months: Creating an AI fast lane for grid connection



Zachary Skidmore
Senior Reporter
- Energy and
Sustainability

DCD gets the inside track on how GridCARE is using AI to find extra capacity across the grid

Getty Images

The University of Stanford has long been a breeding ground for innovation. The university where Google, among others, was founded, remains one of the most fertile grounds for new and exciting companies.

One of the latest to emerge is GridCARE, which formally launched earlier this year following the successful raise of \$13.5 million as part of its seed funding round. The company, born out of Stanford's Sustainability Accelerator, aims to address one of the biggest bottlenecks facing the data center

industry today: time-to-power. To do so, it is leveraging one of the fruits of the sector's unprecedented growth, artificial intelligence.

The company has launched a generative AI platform, which it claims can unlock untapped capacity across the US grid, enabling developers to secure reliable power at a fraction of the time and cost, without having to wait for costly transmission upgrades.

DCD spoke with its cofounder and CEO, Amit Narayan, to learn more about how GridCARE is using AI to make grids more flexible and intelligent in the face

of increasing uncertainty about whether current infrastructure can manage the exponential growth of the data center sector.

Myth Buster

GridCARE was founded on the realization that while the US power grid is almost universally perceived to be heavily constrained, the reality is actually much more nuanced.

According to Narayan, on average, the US grid only operates at about 30 to 40 percent of its maximum utilization capability, meaning that there is

significant untapped capacity that could be unlocked through more intelligent analysis and planning. The founders of GridCARE saw an opportunity brewing to “take the latest advances in generative AI and apply it to improving the power grid,” according to Narayan.

In doing so, the company believes that it can bust a few of the myths of grid constraints, which it says are predominantly the result of outdated assumptions and conservative planning practices.

At present, grid planning typically looks at worst-case scenarios such as multiple outages on high-temperature days and assumes those conditions are persistent across the year. This creates a perception that the grid is constrained when, in reality, constraints are often only present under very narrow and infrequent conditions.

Narayan points to the California grid as the perfect example of this, arguing that it is only constrained during peak hours, with no problems at night or during the winter. GridCARE discovered that there is a significant amount of latent capacity

available on the grid, which is not considered by planners, who fail to take into account operational controls and technologies that are already in place.

Therefore, through the use of generative AI, which considers all the tools and technologies on the grid, GridCARE can simulate and validate its effectiveness under real-world conditions, which it says can significantly accelerate connection timelines and unlock hidden capacity that doesn't typically reveal itself in the utilities analysis.

Increased flexibility

GridCARE is approaching the power issue with flexibility as a central tenet, with AI seen as the perfect tool to unlock greater flexibility across the grid. Its solution works through the use of generative AI and advanced scenario modeling, which identifies geographically and temporally specific constraints and proposes targeted, cost-effective bridging solutions.

The solution is able to analyze hundreds of thousands of possible grid scenarios and pinpoint exactly when and where congestion occurs, allowing for better

“We want to remain neutral, like a TSA PreCheck. We help developers and utilities move faster by creating a trusted “fast lane,” following utility standards, and only challenging assumptions when it’s justified.”

>> Amit Narayan, CEO
GridCARE

planning and understanding of where new capacity can be added to the grid. These include leveraging existing tools like demand response programs, battery storage, and microgrids more effectively to free up extra capacity on the grid.

Through its solution, the company seeks to support utilities that have access to operational tools like battery storage and demand response, but have not factored them into planning. In addition, it works directly with data center developers, from major hyperscalers to AI data center developers, to accelerate time-to-power for infrastructure deployment, both for upgrading existing facilities and identifying new sites with immediate power availability for gigascale AI clusters.

The impact of this can be staggering in terms of reducing the waiting times for a connection, says Narayan, with data centers that usually would have to wait between five to seven years for a connection, seeing time to power slashed to potentially six to twelve months.

“We don't just model constraints—we also map what assets are available to relieve those constraints in real time,” Narayan says. “By combining grid data, asset visibility, and scenario modeling, we’re able to surgically unlock capacity—even in regions where traditional thinking says nothing is possible for the next five to seven years without building new infrastructure.”





According to Narayan, GridCARE's solution has already seen "tremendous interest from data centers across the spectrum." This interest is not only confined to the major hyperscalers, with data center operators focused on inference in a prime position to take advantage of the solution, due to their inherent flexibility, both geographically and temporally. This means that they can be sited in areas with a smaller footprint, which will expedite their time to power.

In addition, the solution offers significant financial benefits for the data centers themselves. This is especially true for AI data center developers, where revenue is greatly impacted by power availability, with the typical estimate being about \$10,000 of lost value per megawatt per day of delay.

Narayan notes that this tends to add up quickly, as for a "100 to 200MW project, you're talking about hundreds of millions or even billions in potential value. Therefore, our technology doesn't just help them find power, it helps them accelerate time to revenue and reduce the need for expensive new builds by better utilizing the assets we already have."

Acting as the middleman

Despite its clear focus on reducing time-

to-power for data center developers, GridCARE has positioned itself as a "trusted, neutral third party" between the data centers and utilities that serve them, says Narayan.

As a result, GridCARE is not only partnering with data center providers but also the utilities themselves. It has already signed partnerships with Portland General

Electric and Pacific Gas & Electric, which view the solution as a means to better utilize their grid assets, as a means to increase overall revenues, and bring down the costs of electricity.

"Collaborating with GridCARE and using advanced planning tools enables Portland General Electric to make more informed and faster decisions in bringing this critical infrastructure online with confidence," said Larry Bekkedahl, SVP of advanced energy delivery at Portland General Electric.

Therefore, like Switzerland, neutrality is embedded in GridCARE's approach, which in turn will allow for greater flexibility to support not only data centers, but any large load seeking an expedited connection to the grid.

"We want to remain neutral, like a TSA PreCheck. We help developers and utilities move faster by creating a trusted "fast lane," following utility standards, and only challenging assumptions when it's justified," says Narayan.

In acting as a middleman, GridCARE hopes to improve communication between the different stakeholders, which Narayan accepts is poor, which has led developers to view utilities as blockers, and utilities to consider developers with trepidation, unsure whether their

"At present, utilities don't report utilization or queue backlogs. Just measuring those would be a great first step. And studies have shown that modest flexibility, such as a battery, demand response, or a virtual power plant, can unlock significant capacity,"

>> Amit Narayan

proposal is serious or speculative. Therefore, in acting as a middleman between the two, both sides are provided greater clarity on when power can likely be delivered, removing the sense of distrust that has become prevalent throughout the industry.

In positioning itself as a middleman, GridCARE has also developed a somewhat unique business model, with the company being paid for every successful transaction that happens between the utility and the developers. Therefore, as Narayan contends, GridCARE has skin in the game, meaning that it is extra incentivized to see results.

Regulatory landscape and US focus

For Narayan, the “secret sauce” of the company is the fact that it works within the existing regulatory regime, meaning that its deployment is not dependent on any major policy change from the Federal Energy Regulatory Commission (FERC) or the North American Electric Reliability Corporation (NERC).

Despite the secret sauce, Narayan argues that greater transparency across the regulatory landscape would still make a significant difference, especially on the side of utilities. Currently, utilities do not typically report utilization or queue backlogs, which significantly hinders the ability to understand how modest flexibility, in the form of a battery or virtual power plant, could unlock significant flexibility across the grid.

As a result, if this flexibility were factored into grid studies, and especially if the public utility commissions required it, it could create much greater urgency within the sector to support solutions such as GridCARE’s.

For the foreseeable future, GridCARE is focusing all its energy on the US market, due to the size of the country’s data center sector and massive load growth projections. The company plans to focus predominantly on regions facing the most acute capacity bottlenecks, including California, Texas, and the Northeast. However, grid constraints are not simply a US issue, but one that is wreaking havoc across the globe.

Therefore, while GridCARE will continue to focus on the US market, it contends that all its solutions are globally

“We don’t just model constraints—we also map what assets are available to relieve those constraints in real time. By combining grid data, asset visibility, and scenario modeling, we’re able to surgically unlock capacity—even in regions where traditional thinking says nothing is possible for the next five to seven years without building new infrastructure,”

>> Amit Narayan

applicable. Narayan revealed that even without a proactive marketing campaign, it has seen interest from across Asia, the Americas, and the Middle East. Therefore, if the company can demonstrate success in the US market, there is the potential for global expansion down the line.

No time for a grid connection

The ultimate goal for GridCARE is to create a system where grid connection time is reduced to zero, effectively

removing the bottleneck of time-to-power.

“We want to go to a point where there is no wait for getting connected to the grid,” says Narayan.

The goal is very ambitious, given the current state of the US grid and the huge projections of load growth from the data center sector alone. To achieve this, Narayan says that GridCARE is seeking to redefine how grid capacity is understood, planned, and utilized, using the current infrastructure more intelligently.

“We’re not just a technology company—we’re helping to shift mindsets. We’re challenging assumptions that no longer reflect today’s reality,” says Narayan.

GridCARE is not the only firm exploring the use of AI in this regard, with several companies emerging in 2025 alone promising significant cuts in grid connection through proprietary AI tools.

Several companies have launched or signed deals seeking to deliver similar reductions in grid connection times through AI. Most notably, Google X’s moonshot project Tapestry, which signed a deal with PJM Interconnection to manage its interconnection queue and automate processes currently completed by grid planners, ultimately seeks to create a model of the grid, akin to Google Maps.

Therefore, we are seemingly on a precipice, like many other sectors, where AI companies are offering solutions that could revolutionize the way we interact with some of the most critical infrastructure. As a result, we could be getting closer to a day where time-to-power no longer represents one of the biggest bottlenecks for data centers on their route-to-market. ●




Getty Images

Go with the flow



Matthew Gooding,
Senior Editor



Are flow batteries the answer to the energy storage needs of data centers?

The soaring energy demands of data centers in the AI era have rapidly become one of the biggest headaches for executives in the industry.

Not only are racks getting denser and more power-hungry than ever before, but they are doing so against the backdrop of an energy grid that is, in many parts of the world, transitioning away from fossil fuels towards more sustainable sources of power.

Renewables, by their nature, are less consistent than fossil fuels when it comes to supplying energy, so battery energy storage systems, better known as BESS, are being delivered at many new data center developments. These systems collect and store energy at times of surplus, meaning it can be redirected to a data center - or back into the wider grid - at times when the wind drops or the sun isn't shining.

But while the benefits of BESS are well established, the type of battery that should sit at the heart of these systems remains up for debate. Cells based on the lithium-ion chemistry currently dominate the market, but for large-scale installations, these can be problematic. Damaged lithium-ion batteries can cause fires or explosions, while geopolitical tensions between China and the US mean getting hold of the materials required to make them in the first place is a challenge.

In this context, businesses and researchers are looking to other chemistries, and flow batteries, a type of cell that doesn't rely on lithium, could prove to be a winning formula if challenges around its adoption can be overcome.

Everything must flow

Flow batteries differ from conventional

cells because they use a liquid electrolyte to store energy, rather than a solid material.

"You have two tanks, one positive and one negative, with the charged storage material dissolved into a liquid," explains Tom Sisto, CEO of XL Batteries, which makes grid-scale flow batteries. "From those tanks it's pumped through the cells, and as it flows over the electrodes, it does the charging and discharging goes back into the tanks."

With a flow battery, you can scale up the size of the storage tanks without needing a corresponding increase in energy, so in theory, they make an ideal storage option for squirreling away excess power. The technology has been around for years, but the liquids used in the electrolyte have traditionally been quite problematic.

"Vanadium is the classic one, which is a highly mature technology that has been scaled for grid use," Sisto says. These batteries use vanadium ions as the charge carrier, and Sisto explains: "Vanadium is expensive and found in geopolitically sensitive regions like Russia and China specifically, and it only dissolves in sulfuric acid. This means the electrolyte

solution of vanadium and sulfuric acid is very caustic and drives a lot of component costs within the device."

XL Batteries' solution is to use an organic compound dissolved in pH-neutral water. Its technology is based on research from Columbia University in New York, where the company's co-founders discovered organic molecules that are stable in both charged and discharged states, making them suitable for incorporating into a battery.

Vendors of flow batteries claim they offer significant advantages over lithium-ion for data center operators looking to install a BESS, the chief one being that the likelihood of water-based cells catching fire is, for obvious reasons, fairly remote.

"The safety aspect is really important for the data center folks," says Giovanni Damato, president of the US division of CMBlu Energy, which manufactures what it calls an organic solid flow battery, using a water-based electrolyte. "I think everyone's concerned that their batteries do not have thermal runaway events in the communities where their data centers are located. Our technology doesn't have that risk."

Thermal runaway can occur when a lithium-ion battery becomes old or damaged, and leads to the temperature of the battery increasing rapidly, running the risk of fires or explosions.

Away from safety, flow batteries could also be more cost-effective and efficient than lithium-ion for long-duration storage, which the US Department of Energy classifies as ten hours or more. "In the eight-to-ten hour storage range, our prices are comparable to lithium-ion right now at low volume," Damato says. "As we scale up and can buy materials in bulk, we're going to see that price drop."

"The other benefit for the type of loads required in data centers is that we can cycle a lot. Our batteries don't degrade based on cycling like lithium-ion does. It's more of a steady calendar life degradation, so you're looking at a 10-15 year lifespan before you have to refurbish the battery."

Shock absorbers

With plenty of potential benefits, it's no surprise data center operators are looking at liquid flow batteries as part of their energy solutions.

"The safety aspect is really important for data center folks,"

>> Giovanni Domato

XL Batteries has agreed to a partnership with Prometheus Hyperscale, which will see batteries deployed at the company's data centers over several phases. In the initial phase, XL will supply and install a 333kW demonstration-scale, standalone Organic Flow Battery at an undisclosed Prometheus site in 2027. Following this, Prometheus plans to acquire a 12.5MW/125MWh commercial-scale system in 2028, with another identical system to follow in 2029.

"We fit a lot of use cases, but for Prometheus specifically, we'll be used as a 'shock absorber' for large compute power swings," Sisto explains, referring to the wildly varying power draw of GPUs, which DCD reported on in April. "So as they ramp power up and down very quickly, we'll be there as a buffer."

Sisto believes the technology will be useful, not only for storing energy from intermittent sources of energy like renewables, but also for "flat" energy generation technologies like small modular nuclear reactors, which are expected to be deployed in data centers over the coming decades.

"You don't shift nuclear power generation quickly," he says. "So storage allows you to match the variability of demand with a flat line of generation."

CMBlu's technology is being deployed at a DataHall data center in Saale, Germany, which is set to come online in 2027/28. It will initially feature a 4MW battery, with the



aim of expanding to 50MW.

"We've got a few customers in Europe that we're working with because they're facing specific local constraints around electricity supply and pricing, and want a battery on site," Domato says.

Challenges ahead

While flow batteries show great promise, issues around scaling the technology remain.

Last year, Australian vendor Redflow Batteries went into voluntary administration, having failed to secure investment for its flow battery, which used a liquid electrolyte combined with zinc-bromide in what Redflow claimed to be an environmentally-friendly solution.

Despite receiving multiple government grants, and landing a contract to supply 15.4MWh of batteries for a microgrid in California, it said it had been "unable to attract the required equity support" to continue, and administrators were unable to find a buyer, demonstrating some investor reticence around the potential of the batteries.

With lithium-ion being such a well-proven technology, Domato admits flow batteries still have a way to go before they are used widely in data centers and beyond.

"Lithium-ion has taken 60 years to get where it is today," he says. "We're getting close to a full commercial roll-out [for flow batteries], but it's going to take a lot of patient capital to keep things moving. We're in that valley between pre-commercial and commercial, and ramping up is our key challenge right now."

XL Batteries' Sisto is confident flow batteries have a role to play alongside other storage technologies as data centers navigate the energy transition. "The global energy market is one of the largest markets in existence," he says. "The numbers we're talking about are so astronomical that they're almost incomprehensible."

"If we're going to modernize what is the most complex machine in the world, I think it will take everybody, and I am hopeful that the industry and the markets will move together towards a new era of energy generation, storage, and consumption." ●



End-to-End Power Monitoring. Accuracy you can trust.

PowerLogic™ metering delivers real-time insights and advanced power quality analysis to detect issues early, reduce risk, and ensure reliable, high-performance operations.

PowerLogic ION9000, PM8000 and HDPM6000



[Download the eBrochure](#)



Zachary Skidmore
Senior Reporter, Energy
and Sustainability

Drilling for data: Can geothermal energy meet hyperscale ambitions?

Meta and others have thrown their backing behind experimental geothermal projects as AI energy demand continues to rise



The sound of drilling echoes in the New Mexico sunshine. You'd be mistaken for thinking this is just another instance of fossil fuel exploration; however, the large hole being bored into the ground serves a different purpose: to generate energy through heat resonating from the Earth itself.

Often, we forget that the deeper you go into the earth, the hotter it tends to get. Humans' utilization of this heat, better known as geothermal energy, dates back to ancient times, with evidence of its use in bathing, cooking, and washing, most commonly via underground hot springs.

But it's only since the early 20th century that humankind has been able to harness the Earth's heat to generate

electricity, with the first geothermal power plant built in Larderello, Italy, in 1911.

Today, we stand on the threshold of a new geothermal revolution, with innovative techniques, known as Enhanced Geothermal Systems (EGS), offering a new and innovative way of harnessing geothermal energy at much higher capacity factors and across a significantly larger footprint.

Its potential has not gone unnoticed among the hyperscalers, with Google, Microsoft, and Meta all signing agreements with geothermal developers in the last two years alone.

But can geothermal energy become a key source of power for the data center sector, both in the US and further afield?

Heating up

Geothermal isn't a particularly new technology. Shallow conventional techniques, utilizing heat pumps and borehole heat exchangers, have been in common use since the 1970s.

Conventional wells are typically up to 400 meters (1,312 feet) deep and used mainly to heat and cool buildings. While we have seen a growth in deeper geothermal electricity generation projects, which reach depths exceeding one kilometer (0.62 miles), these wells are very location-specific, found in regions with high-temperature underground resources, such as the US, Indonesia, and Kenya. As a result, geothermal energy currently accounts for less than one percent of the global energy mix.

EGS, unlike conventional methods, has the potential to boost that figure significantly, by tapping into “hot dry rock,” which is more abundant than natural hydrothermal reservoirs and offers a substantially higher generation potential.

EGS was first explored as a concept in the 1970s at the Los Alamos National Laboratory in New Mexico. Inspired by offshore oil drilling techniques, the scientists tested the possibility of forming artificial geothermal reservoirs by fracturing deep, hot, impermeable rock. The project proved a success, producing 10MW of energy by 1985. But cost and complexity issues saw it canned.

Interest in EGS was revived in the 2000s, with considerable advances in hydraulic fracturing and other drilling techniques leading to a resurgence of investment. A 2006 MIT report stated that: “Geothermal energy from EGS represents a large, indigenous resource that can provide base-load electric power and heat at a level that can have a major impact on the United States.”

Today, the International Energy Agency (IEA) estimates that in the US, at least 7TW of capacity is accessible at depths less than 5km (3.1 miles), and more than 70TW is accessible across all depths. Most of the potential capacity is located in the Western US, specifically in California, Nevada, Utah, Oregon, and New Mexico.

In comparison, the “technical potential estimates for hydrothermal energy are about 25GW,” says Ben King, director at Rhodium Group’s energy and climate practice. “With EGS, you’re on the order of terawatts available across much wider areas of the country.”

This vast potential has led think tanks such as the Rhodium Group to project that, if scaled effectively, EGS systems could ultimately supply nearly two-thirds of new data center demand by 2030.

Energy in abundance

Meta has been a particularly enthusiastic backer of EGS.

Facebook’s parent company signed deals with EGS firms Sage Geosystems and XGS over the past 18 months. The two projects differ significantly in terms of location and technology, providing a glimpse into how such systems actually work.

In August of last year, Sage Geosystems

“The technical potential estimates for hydrothermal are about 25GW. With EGS, you’re on the order of terawatts available across much wider areas of the country,”

*>> Ben King,
Rhodium Group*

became the first EGS company to partner with Meta for 150MW of power from its project located “east of the Rockies.” Founded in Houston by a former oil and gas executive, Cindy Taff, the company leans heavily on techniques developed in that industry.

At the heart of Sage’s offering is what it has dubbed a “pressure geothermal system,” which combines heat and mechanical energy stored deep underground.

“We call it pressure geothermal because we’re using not just the temperature underground - we’re using the pressure element as well. There’s thermal energy, and there’s mechanical energy. When you combine them, you can increase the net output of your system,” explains Jason Peart, general manager of strategy and development at Sage Geosystems.

Sage drills paired wells, reaching depths of up to 20,000 feet (6.1 meters) into hot rock at temperatures of 180°C (356°F) or higher. In each well, it creates an artificial “lung-like” reservoir that is filled with water. The rock’s natural elasticity allows it to flex and contract, storing pressure in addition to heat. One well produces hot, pressurized water to a surface power plant while the other recharges and reheats, with the two switching roles daily.

The system is modular by design, which creates incredible potential for scaling, says Peart. “Each well pair has a net output - 3MW in some places, 8MW in others - and you just keep adding additional well pairs for greater capacities. If you want 100MW, 500MW, even a gigawatt, you put more pairs on a pad.”

The ability to scale doesn’t only boost output, Peart argues, but also makes it more economical.

“One well pair isn’t particularly cost-efficient,” says Peart. “It carries all the costs of mobilizing equipment and building the facility. But when you start putting eight, ten, 12 wells on a pad, you get drilling efficiencies, operational efficiencies, and economies of scale that change the cost dynamics.”

In order to change the cost dynamics, Sage needs to demonstrate that its technology actually works. It currently has several projects in its pipeline, mostly located in Texas. The Meta project has no fixed location and is expected to be developed via a phased approach, with phase one of four-8MW completed in 2027 and phase two of 150MW completed in 2029.

Closed loop

In its second deal, signed in June, Meta partnered with XGS on its planned project in New Mexico, a state the vendor’s CEO Josh Prueher calls one of the “world’s best geothermal resources.”

While the region offers an abundance of hot rock, its arid conditions make water use a challenge. To address this, XGS has transitioned away from traditional water-based geothermal energy and is instead pioneering a fully closed-loop geothermal system.

XGS starts by drilling a single well into rock hotter than 200°C (392°F) and lowering a steel casing to the bottom. It then injects a ‘thermal reach enhancement’ material to draw heat from the surrounding rock. Inside the casing, an insulated tube completes a tube-within-a-tube system, carrying the captured heat to the surface.

The design ensures that water never comes into contact with rock, thereby eliminating water loss and contamination. Additionally, according to Prueher, the system is pressurized, sterile, and chemically-controlled, ensuring predictable and stable thermal flows over decades. This lowers operational risk and costs, Prueher says, as unpredictability can often lead to a 40-50 percent drop in the effective capacity of geothermal plants.

“Banks and independent engineers want to see steady, predictable flow rates over 20 or 30 years,” he says.

"Conventional systems just haven't been able to give them that because nature is variable. That's why geothermal has been stuck at about 4GW in the US for decades."

In providing this, XGS views itself as competing on the same terms as other baseload power sources, rather than renewables. "Our value proposition is: we're clean like renewables, but we're firm like natural gas - and we can site close to load. That's why data center operators are taking this seriously," Prueher contends.

The project is expected to launch by 2030 and supply 150MW into the local grid, which will in turn supply data centers in the region. The company intends to develop several projects in the region, ranging from 5MW to more than 500MW.

Like Sage, the company is looking to lean heavily on the oil and gas sector to utilize the drilling technology required to access the hot rock deposits. "We're not waiting on some moonshot," says Prueher. "That's why we're confident we can go from first-of-kind to gigawatt-scale projects in just a few years."

Strategic bets


So what has made Meta so confident in the technology that it has signed not one but two deals within the sector? For John DeAngelis, the company's head of clean technology innovation, the agreements are "strategic bets."

"Those bets are designed to help technologies and companies scale, to prove their technical feasibility at scale, and to drive down costs in an accelerated way," he says.

As of yet, only Fervo Energy, which became the first EGS firm to sign a deal with a data center company back in 2022 when it inked an agreement with Google, has successfully initiated a demonstration project. In 2023, the company reported positive results from a 30-day well test conducted at Project Red, its full-scale EGS pilot project in northern Nevada.

Therefore, according to DeAngelis, the next few years will prove crucial in determining the future of EGS, demonstrating whether it can perform effectively and at scale. "Success for us would be seeing hundreds of megawatts - if not gigawatts - of geothermal deployed by 2030," says DeAngelis.

Supporting the sector in commanding the cost curve will be a crucial factor, as driving economies of scale will enable the



technologies to develop more effectively and compete with other sources of energy available on the market.

The support provided by the data center sector could prove an effective tool to drive these scales, says Annick Adjei, senior research analyst for subsurface at Wood Mackenzie. "The support from companies like Meta, Google, and Microsoft goes far beyond traditional power purchase agreements," she says. "These are real partnerships helping geothermal startups commercialize their technologies and bring projects to life."

By becoming "foundational" backers, the companies can drive greater investor confidence in the sector, argues Adjei. And it's showing, Wood Mackenzie has tracked more than \$2 billion in capital

raised by next-generation geothermal companies since 2019, with the majority coming from private investors.

However, even with all the money in the world, the technology itself needs to prove that it can provide reliable power. Therefore, for Adjei, "to truly understand the full potential of enhanced geothermal systems as a viable technology, we need far more demonstration projects beyond [Fervo's] Cape Station in Utah."

In front of or behind the grid

If EGS is successful in becoming a cost-effective and scalable power source, the question then becomes how it will serve the data center market. Currently, hyperscalers are focusing on integrating the new geothermal capacity into the grid rather than supplying power to data centers behind the meter.

"Deploying energy resources on the grid can be not only the fastest but also the most reliable option," says DeAngelis.

While Meta and Google have yet to commit to off-site power, siting data centers close to or adjacent to geothermal wells is a clear possibility. "We believe the next big opportunity is true colocation - data centers built adjacent to geothermal wells. It eliminates transmission risk and makes permitting cleaner and potentially faster," Adjei notes.

Due to geothermal's baseload profile, it can be deployed in an "island grid" setup, bypassing interconnection delays and avoiding public backlash over huge energy consumption.

Geothermal developers emphasize the added benefits and potential as an off-grid solution for data centers. "For a hyperscale data center, we can give you firm, clean, 24/7 power right behind the meter, plus you get cooling load support and clean water production as part of the same footprint," says Prueher. "That's not just a power solution - it's an infrastructure package."

As a result, Peart predicts that in the future, data centers may construct their facilities near geothermal plants for direct supply, following a model similar to natural gas. However, such an approach relies on EGS scaling to gigawatt levels and demonstrating its reliability.

Interest in geothermal as a behind-the-meter option is not limited to the US, as Simon Westerlund, investment manager at Baseload Capital, notes. "We



are already seeing data center players looking to be located close to geothermal power plants in places like Indonesia and the Philippines," he says.

Subsequently, we could end up seeing geothermal energy act in a similar way to natural gas in the US, in both supplying firm grid power and providing an expedited route to market for data center operators who are willing and able to build their facilities adjacent to a geothermal well.

Barriers and limitations

Despite the fervor and heavy investment into EGS and its apparent applicability to the data center sector, concerns remain, one of the most notable being the sheer cost of drilling. Studies from Stanford and the US National Renewable Energy Laboratory indicate that drilling alone can account for between 30 and 57 percent of a project's upfront capital costs, depending on the system's design.

For operators of data centers, which demand both baseload reliability and tightly managed electricity prices, such figures represent a major barrier. Even as new techniques promise to reduce these costs, the challenge remains. "If wells can be drilled as quickly and cost-effectively as oil and gas wells, that will speed deployment and reduce costs," says Rhodium Group's King.

Cost concerns remain only a part of the story, with the act of drilling itself coming with significant risks. Geothermal wells, especially deep ones, are vulnerable to technical setbacks, including circulation losses, equipment failures, and mineral scaling, which can compromise flow.

A report from the Clean Air Task Force notes that such problems can delay projects, inflate budgets, or even force developers to abandon wells entirely. Additionally, the need for specialized drilling equipment to access the hot, dry rock poses a risk that the supply chain could become constrained, potentially

leading to further delays.

The regulatory landscape compounds the difficulty. "The main regulatory hurdle is the permitting process to drill wells, which is handled mostly at the state level, with some EPA involvement for water permits," King argues. Lengthy approval processes can add years to timelines, a prospect that is unpalatable for hyperscale data center operators, whose expansion plans often run on tight schedules.

The tide may be turning on this, however, with policymakers across the aisle in the US voicing support for the sector, exemplified by US Energy Secretary Chris Wright's ardent backing of the technology. This is something that hasn't gone unnoticed by Meta, with DeAngelis noting "we've seen a groundswell of support for geothermal."

Finally, there is the question of cost competitiveness. A study from Environmental Research showed that even halving geothermal costs by 2050 would leave the technology struggling to match the cost-effectiveness of solar and storage. Only if costs fall by more than 70 percent would geothermal emerge as the most affordable carbon-free option.

Is conventional still a possibility?

While the US leads the way in data center-linked geothermal deals, there have been several other instances further afield. Unlike in the US market, the agreements have focused on conventional geothermal energy rather than EGS.

This is best exhibited by Google, which earlier this year signed a 10MW Power Purchase Agreement with Baseload Capital in Taiwan. Unlike US deals that center on experimental EGS, the Google arrangement will rely on conventional, shallow wells.

Baseload's Westerlund explains the reasoning: "We are mainly taking geological risk, not technology risk," he says. "So we're sticking with conventional

technologies for our first batch of projects."

According to Westerlund, the decision reflected the realities of the Taiwanese energy market. Though the island has significant geothermal potential - between 30 to 60GW according to some estimates - it lacks a mature oil and gas industry to supply rigs and expertise at the rate required for EGS.

As Westerlund puts it: "The US is faster to move than Taiwan because the oil and gas ecosystem already exists. Rigs are available, and it's easier to mobilize resources."

As a result, Baseload saw conventional systems as its best bet, and views the Google deal, which included a direct investment in the company, as a means to expedite the development process of its projects in Taiwan. Consequently, Baseload expects to reach commercial operations at its first sites in the country in 2029.

In addition to the PPA, the deal is expected to foster the growth of Baseload's geothermal footprint across the Pacific, with Westerlund arguing that it's not just the US that is in a position to take advantage of geothermal energy. "Japan, Indonesia, and the Philippines are pushing geothermal again - they are restarting their industries with government support," he says.

These projects are very likely to be conventional-based, as they face similar challenges to Taiwan in accessing drilling equipment and expertise to deploy EGS systems at this time.

However, if successful, they could not only support the growth of conventional geothermal systems but also provide a springboard for the deployment of EGS. This is reflected in Baseload's technologically agnostic approach to geothermal, which means they will be open to exploring EGS if its level of technological risk is reduced, potentially expanding its reach across the Pacific and beyond.

Consequently, conventional geothermal energy clearly has a role to play, particularly in regions with accessible shallow resources and established drilling capacity. Yet its reach is limited. For the data center sector to truly harness geothermal at scale, the industry will likely need to move beyond these pockets of opportunity and embrace the wider potential offered by EGS, in the hope that its "strategic bet" will pay out. ■



The power behind your data.

End-to-end grid connection solutions across the power landscape.



When it comes to electricity for data centres, playing it safe is the only option. From assessing a site's feasibility and securing a connection, to design, build, ownership and operation - Eclipse Power Group have the experience and expertise to get you connected.

With simplified innovation in every step, future ready grid solutions and asset value paid on adoption - your pre-2030 connection journey could be simpler than you think...



Come and see our team of friendly experts
Stand 119, DCD Connect 2025

**Electricity Connections.
Simplified.**

Going local in Europe; the rise of the regional Edge



Dan Swinhoe
Managing Editor

Data center veteran Adriaan Oosthoek talks about his new firm, Portus



In the rush to serve the hyperscale operators, the needs of local enterprises, especially in Europe, risk being forgotten.

The big colocation providers are scurrying to deliver massive campuses for the traditional cloud providers as well as the raft of new AI neoclouds. But a new wave of Euro-centric operators see opportunity in serving local customers in markets away from the traditional FLAP-D locales with smaller facilities.

Portus Data Centers is one such European colocation operator looking to build out a platform focused on serving the local needs of local users away from the traditional hyperscale markets.

"Big is relative," says Adriaan Oosthoek, chairman of Portus. "A 5.5MW data center in Munich for Digital Realty or whoever, it's completely subscale. For NTT, it is too small. But it's big enough for us."

A new local European operator

Portus was formed in December 2023 as a roll-up of investment firm Arcus Infrastructure Partners' data center acquisitions.

December 2022 and January 2023 saw Arcus acquire Luxembourg operator European Data Hub (EDH) and Munich's SDC SpaceNet DataCenter in Germany. After the acquisition of Internet Port

Hamburg (IPHH) the following December, the three regional facilities were grouped together under the new Portus name with Oosthoek as chairman. The three companies were acquired by Arcus's European Infrastructure Fund 3 SCSp (AEIF3); terms weren't shared for any of the deals, but they gave Portus 5.1MW of data center space.

"The premise here is a buy-and-build strategy in order to build a regional colocation platform," Oosthoek tells us during our chat in March 2025. "The plan is to add one or two more diverse locations, probably in Germany, over the next 12-18 months."

Oosthoek was previously SVP of operations, construction, and IT at Interxion until 2022. He has also held roles at Colt Technology Services, TeleCity, Redbus Interhouse, Teles, and Racal Datacom.

Oosthoek says that, as well as acquiring quality assets, the company has been keen to lean into the benefits of buying experienced local teams that know their facilities and customers inside out.

"There is 25 years of history in Hamburg, and about 15 years in Luxembourg. Portus is an interesting mix of new startup as well as a lot of history and people that have been with the company for many years."

"Even in Germany, when you speak the language and the accent in Bavaria, that helps. We have a local guy as general manager in Munich for that reason. He knows the market, he knows the culture, he knows the people, he knows the infrastructure."

Buy, build, expand

Portus is primarily targeting regional cities in Germany, and will most likely be avoiding Frankfurt, the country's hyperscale heartland. While the company has a preference to buy established data centers, Portus isn't afraid of greenfield development.

"We have a list of cities where we want to be, and we are exploring greenfields and acquisition opportunities," Oosthoek says. "When we acquire assets, we use that as our hook into that market and then find expansion opportunities alongside the existing business."

He notes greenfield developments would see Portus targeting facilities with 4-10MW of IT capacity.

"Greenfield is a little bit more difficult in a new market. But if there's no opportunities to acquire, then greenfield is the way to go. We're working on both."

As well as expanding to new sites, Portus is building out the footprint it already has. The company has announced

planned expansions at all three acquired sites, including the addition of 2.2MW in Luxembourg and 12.8MW across two buildings in Hamburg.

An underserved market

Portus is being set up on the premise that there's a sizeable niche to serve companies that need colocation capacity in Tier 2 markets, but want something with more scale than a local provider with one location can provide.

"It's really clear that in Europe, the big guys are focusing on their big customers," says Oosthoek. "The likes of Digital Realty, Equinix, Vantage, and all these guys are running 100 miles an hour just to serve a handful of big customers. It takes all of their energy, all of their capital, all of their mindshare. They really don't have any time left or money left to develop smaller locations."

"Those guys are not investing currently in regional locations. And where they are present in regional locations, you even see some of these guys moving out, divesting smaller sites, because they want to focus their very large campuses to serve the large customers."

One driver is sovereignty. European customers can often be wary of putting sensitive data in facilities owned or operated by US company, so having a European-owned local provider makes an attractive alternative. Portus already serves a number of European institutions out of its Luxembourg site.

"You see the big cloud providers delivering and developing sovereign cloud solutions," says Oosthoek. "But there's a sentiment that customers simply don't trust the sovereignty of those solutions. They want to see things that are proven to the European or German, and we were working on a number of opportunities right now that are going in that direction."

Beyond the demand for sovereign capacity, Oosthoek says he is still seeing local regional demand, especially around local governments, hybrid cloud deployments, or companies repatriating workloads back from the public cloud.

"If an enterprise in Munich repatriates workloads from the public cloud, which is probably sitting in Frankfurt, there is no reason for them to put it in a data center in Frankfurt," he says. "There's more reason for them to put it in a data center near to them that's local in Munich. That's starting



Adriaan
Oosthoek

to appear more and more as real demand."

On AI, Oosthoek says the company is seeing GPU deployments from smaller enterprises, often academic research, industrial robotics, of healthcare, but they tend to be around 15-20kW rack, which can be air cooled.

He expects lower-density inference workloads to be a major driver for regional AI deployments in the future.

"There's probably 15-20 companies out there deploying infrastructure all across the world, and more and more in regional locations," he notes. "These are much more either AI inference-type workloads, which we're starting to see a little bit, CDN workloads, or more network-like workloads to serve local customers in the region."

The company hasn't looked to retrofit existing sites for liquid yet, but is designing its expansions and new-builds to cater for both liquid and air with densities up to 50-60kW per rack – even if it isn't seeing requests for that yet.

"The name of the game in colocation is making sure you design such that you remain flexible," Oosthoek says.

Going local internationally?

Founded in 2009, Arcus is an independent infrastructure fund manager with a focus on the European mid-market, with some €9.6 billion (\$11.2bn) in assets under management. Beyond Portus, it has previously invested in tower companies TDF and Shere Group, fiber firms Swiss4Net, E-Fiber, and FixMap, and numerous other firms in the industrial and energy sectors.

The company isn't the only investment firm looking to build out new 'regional Edge' operators targeting Tier 2 markets across Europe via a build and buy strategy, showing investors believe there is a real

demand for local regional data center capacity. The likes of NLighten, Etix, NorthC, and Netherlands-based Switch Datacenters are among those targeting smaller European markets.

Portus is focused on serving Germany and adjacent countries. On going further afield, Oosthoek says Arcus is firmly focused on Europe and isn't looking to expand Portus into Asia, Africa, or the Americas. He notes, however, that Arcus could set up new data center platforms for other parts of Europe that could work together with Portus, although there are no firm plans to do so yet.

The company is hoping that by serving local enterprises via a regional platform covering a few markets with in-market expertise it can offer a niche that isn't currently well-served.

"The competition is local players, and this is where we can excel," Oosthoek says. "We're planning to build a bigger company with a number of locations. So with us they also got more ability to go elsewhere if they wanted to. Not just one or two sites, but maybe ending up with 15 or 20." ■

AI IMPLOSION?

While Portus isn't directly chasing the large-scale AI contracts, Oosthoek isn't particularly worried about any kind of implosion of the data center industry because of any AI bubble.

He says that we've already seen enterprises shift en masse from on-premise to cloud deployments, and all that happened was those legacy facilities were repositioned to new customers.

"It hasn't caused any real problems because it's been gradual," he says. "Sentiments around it that may turn on its head, because that's what sentiment does, but leases and contracts will always be staggered."

The only facilities he thinks might face trouble are the early hyperscale data center that might be deemed surplus to requirements if the large players need to rationalize capacity and prefer to stick with the newer and higher-density sites. ■

The Arc of Trout

CleanArc Data Centers
CEO on building a digital
infrastructure empire. Again.



Matthew Gooding
Senior Editor



Jim Trout's CV could easily pass for a directory of the data center industry's biggest names.

From CoreSite to Vantage via Digital Realty, the CleanArc Data Centers CEO has been at the forefront of digital infrastructure development for 25 years, and after a spell working as an investor, is now back in the game with CleanArc, which is developing hyperscale facilities, initially in Virginia.

His enthusiasm seems undimmed when he catches up with *DCD* in the wake of CleanArc taking funding from Nuveen, the investment manager for the TIAA pension fund. The size of the investment has not been made public, but given that Nuveen has \$1.3 trillion of assets under management, it's likely Trout and CleanArc have a fair bit of cash burning a hole in their corporate back pocket.

"My view is we source a lot of solar and wind, then add batteries as an option and nuclear as an option"

But, having left Vantage in 2013, did he ever think he'd find himself building another data center company? "I was working on water conservation technologies, making angel investments, and buying up real estate for my own portfolio," he says. "But in 2018, I was looking at the market and thinking 'this is just too crazy.'"

"We're at the next stage of a quantum leap of adoption, a doubling or tripling

of capacity needs. I've always loved opportunities where you can align with your customers, and you can collaborate and have a planning process with them. That's what we're doing with some of the biggest hyperscalers."

Trout makes a splash

Trout started his commercial real estate career with Prologis, before taking a role with investment fund Carlyle Group, helping it "to develop a business plan for its West Coast assets."

These assets would become CoreSite, the data center and telecoms infrastructure firm Trout helped found in 2001, serving as CEO for two years when it was known as CRG West and operating out of one of the original carrier hotels, One Wilshire in Los Angeles. "It was very network-centric, and difficult to scale in certain areas," he recalls. "I said to them,

'this is going to be a great business, but we need to get into a larger data center experience to be able to grow with corporate America.'

Having set CoreSite on the path to growth, in 2003 he moved over to the fledgling Digital Realty to "run the operations" of a company that would go on to be one of the largest investors and builders for digital infrastructure. "I had portfolio management and technical operations, and grew that team from eight to 500," he says. "We went public and expanded from there, but the problem with the model at the time was that there was limited modularization, so it was difficult to scale it to true hyperscale."

Trout's desire to build big was matched by another private equity firm, Silverlake, which came calling in 2009, asking him to launch a new business, Vantage Data Centers, initially in Silicon Valley, where it acquired a former Intel data center. "There was an opportunity with this Intel campus redevelopment that I saw and thought 'wow, we can really get high voltage power close to deployment,'" he says.

With what was, for the time, a massive 37MW of power available at the site, it's no surprise Trout's eyes lit up at the prospect. Vantage launched its data center in Santa Clara in 2011, the first of what would go on to be a network of large wholesale colo sites around the world. "We got Vantage to a point where it was monetizable into DigitalBridge's portfolio [DigitalBridge acquired the company in 2017], I'm super proud of what we achieved," Trout says.

Clean growth

Now his focus is on CleanArc, the company launched in 2023 by power investor 547 Energy. It has since taken PE funding from Snowhawk, which became the firm's biggest shareholder in the process, as well as Townsend Group and, most recently, Nuveen.

All this capital will be used to fund the build-out of hyperscale data centers fit for the AI age. CleanArc says it will deliver these data centers to its clients with renewable energy matching agreements already in place, ensuring they fit with the hyperscalers' net-zero ambitions. The company launched an RFP for clean power partners in 2023, receiving expressions of interest from projects totalling 3GW.

How many of these will come to

fruition and over what time period remains to be seen, but it has secured sufficient power to deliver its first campus, in Virginia's Caroline County. This will eventually offer 600MW at full build-out, with 200MW due to be online by Q2 2026, and a further 300MW following by mid-2027.

Who will be moving in has yet to be announced, but Trout says the company is working with two hyperscale clients interested in doing deals for CleanArc infrastructure.

"We have bi-weekly calls and we're helping them through the whole supply chain process," he says. "We're integrating as much as we can in terms of power, land, and then some level of infrastructure, and letting them

looking to other US markets and different products.

"We are working on multiple other sites in other Tier One markets," he says. "We have several under contract, but there's a process to get the power fully fleshed out. What's exciting is that, after the recent investments, we've been able to get into a more aggressive mode."

CleanArc is also investigating other types of data centers, catering to enterprise users, something Trout did to good effect at Vantage, bagging deals with the likes of Boeing.

"We're looking at things that would be more Edge and corporate AI-driven," he says. "We're in active discussions with Fortune 100 companies about what that could look like. They don't have the size

"In 2018 I was looking at the market and thinking 'this is just too crazy'"



decide. Every company's supply chain management is different, so we're being flexible and adjusting while we do these initial deals."

Though not naming the hyperscalers involved, he says CleanArc has carried out technical due diligence - a process that puts the company and its capabilities under the microscope - with both. "Now we need to show them we have the ability to respond to different CPU and GPU designations, not only those that are on the market now but also Nvidia's future designs," he says. "We need to show technological flexibility."

It will be up to the clients themselves to decide what mix of renewable energy offsets the emissions from their servers, Trout adds. "My view is we source a lot of solar and wind, then add batteries as an option and nuclear as an option - it's up to the hyperscalers whether they consider [nuclear] renewable from their perspective," he says.

A changing landscape

Beyond Virginia, Trout says CleanArc is

of the hyperscalers, but we're still talking about 25-50MW facilities, which are big enough for us to work with and can help us diversify."

The data center industry has changed immeasurably since Trout took his first steps with CoreSite, but he says the things that drive him are unchanged. "The excitement of the job for me has always been building small teams of young people and building capability at a platform level," he says. "I really like going in from the ground level, and being the founder of an idea and a strategy that has legs and is scalable."

But he admits he didn't anticipate the magnitude of the sector's growth back when he was starting out in the early 2000s.

"I didn't see it getting this big, this fast," he says. "But it's still a really challenging environment with long sales cycles, with very sophisticated clients that aren't afraid to throw their weight around. But now I can show them my team, and my investors, and that puts us on more of an equal footing."



**Power
Generation**

CONFIDENCE ON

Industry-leading power technology.
Guaranteed 99.999%¹ reliability.
Single point-of-contact accountability.
Global expertise on tap.
Whole life cycle support.

We know what data centre and mission critical facility operators demand. So we've developed even more flexible options with Centum™ Force, our game-changing series of containerized generator sets.



The past, present, and future of Johor



Jason Ma
Junior Reporter

Making sense of Southeast Asia's
fastest growing market

Johor Bahru-Singapore
Causeway - Credit: Getty Images



Princeton Digital Group's 52MW, JH1 data center
Credit: Princeton Digital Group

Johor, the southernmost state on the Malaysian peninsula, is historically associated with rubber and oil palm plantations. In recent years, it has also become one of the world's fastest-growing data center markets.

It helps to get a sense of the facts and figures. Real estate consultancy CBRE says Malaysia possesses the biggest pipeline of data center developments in Southeast Asia. Cushman & Wakefield concurs. And according to a July 2025 report written by another data center consultancy DC Byte, Johor currently offers around 487MW in live capacity; has 324MW worth of projects under construction and 1.4GW worth of committed projects – referring to developments which have secured all necessary elements like energy and land but have not begun construction – and, most stunningly, has 3.4GW of projects in the early stages, which refers to projects which have secured some, but not all, of its required elements.

Compare this to Singapore, which is separated from Johor's state capital, Johor Bahru, by a short 1km causeway. The city-state continues to enjoy its traditional position as the focal point of Southeast Asia's data center industry. Its role as an Asian financial center has produced a robust digital infrastructure market built to support enterprise,

"Customers generally like to put their resources in multiple locations. It's a concentration of risk kind of thing,"

>> Kang Soon Chen

financial services, and cloud workloads, and as a result, offers 1072MW of live capacity as of January - approximately double what Johor currently has to offer.

But the total amount of capacity under construction, both in terms of committed capacity and early stage supply, is 445MW - a far cry from Johor's 1.4GW - as the effects of an ongoing data center moratorium continue to be felt.

What we can gather from this is that Johor is a growth market. One might not be impressed by looking at the quantity of operational data centers – Knight Frank estimated that by the end of 2024, Johor only possessed 12, and Baxtel, which ostensibly provides a live database, estimates 16 – but one would certainly be cowed by what is to come. Compare the Johor of July with the Johor of January. A glance at DC Byte's reports show that within the six-month time span, Johor added an extra 90MW

of live capacity, started building an extra 118MW worth of projects, and gained an extra 1.6GW worth of committed and early stage projects.

A short history of data centers in Johor (and Singapore)

Singapore has long been the predominant data center hub in the region. But this put a strain on the city state's power grid, leading the government to issue a moratorium on new data center developments in 2019.

Though no official statement or legislation was passed to that effect, those working in the industry quickly got the message after it became impossible to secure permits for new facilities. No new construction projects were approved between 2019 to 2022, and the moratorium was only officially acknowledged in late 2020 in the form of a pithy joint statement issued by several government agencies to local broadsheet *The Business Times*.

Up until this point, Malaysia's largest data center market was located in an area called the Klang Valley which encompasses Kuala Lumpur, the country's capital, and Cyberjaya, a planned tech city. According to real estate consultancy CBRE, this market has historically supported a wide range

of cloud-based services, ranging from enterprise digital platforms to banking, finance, and government systems, and took up a lot of the slack created by the Singapore moratorium.

All the American hyperscalers therefore have a presence in the region. The usual crowd – AWS, Microsoft, Oracle – are rounded out by Chinese operators like ByteDance and colocation providers such as AirTrunk, EdgeConneX, and DayOne. But the biggest player so far is Bridge Data Centers. A subsidiary of Bain Capital-backed ChinData, it operated 32 percent of built IT capacity in Malaysia as of December 2024. DayOne operated 29 percent and AirTrunk operated 13 percent. As is characteristic of all data center businesses, they often lease out one another's services. DayOne in particular has reaped great rewards, partnering with Oracle to service ByteDance.

Demand for these traditional workloads are still growing but the overflow of demand from Singapore has had a new focus since ChatGPT debuted in November 2022, sparking a new wave of interest in AI.

Indeed, AI processing has become the key driver for new capacity. The demands of large-scale model training – accompanied by much larger electricity bills than traditional workloads – and inference mean that facilities are being designed specifically for AI.

The center of gravity for data centers in Malaysia has therefore shifted down south to Johor. Within the state, it is the area around Johor Bahru, its southernmost city, that has ballooned. According to Knight Frank's 2024 end of year report, Johor contains about 80 percent of Malaysia's IT capacity, and much of that is within Johor Bahru. The Klang Valley still possesses a greater quantity of data centers, but the difference in workloads means that 79 percent of its data centers offer below 20MW of capacity. Compare this to Johor, where 53 percent of its data centers offer more than 10MW.

But the question of why the data center market entered Johor with such vigor is separate to the question of why it chose to stay. After all, Singapore started to remove its moratorium in 2022, and yet DC Byte and others argue that

"...investments in data centers should not be pursued unless they bring tangible added value to the rakyat [meaning ordinary people], such as high-paying job opportunities and knowledge transfer,"

>> President Anwar Ibrahim

demand will probably remain stronger in Johor than in Singapore for the time being.

So why stay?

The beginnings of an answer start with Johor's natural endowments. It possesses abundant land – just under 20,000 km², which is around 26 times Singapore's measly 735.7 km² – and offers water and energy at much cheaper prices.

The Malaysian government has also instituted policies to make Johor more amenable to new and existing data center operators.

For instance, the government was quick to introduce tax incentives. In 2022, it announced the Digital Ecosystem Acceleration tax incentive scheme, offering new and existing data center operators a tiered investment tax allowance, a special tax rate on statutory income, and other benefits provided they meet a few requirements. This includes the hiring of a certain number of Malaysian employees, the guarantee of a monthly basic salary of MYR 5,000 (around \$1,200), and the adoption of a 'green technology.' According to data from the Malaysian government, up until March 2025, a total of 21 data center projects had been approved under the scheme, drawing in a total investment of RM 113.8bn (\$24.4bn) since its inception.

It has also made it easier for data center developers to acquire a grid connection. In August 2023, the sole distributor of electricity in Peninsular Malaysia, Tenaga Nasional Berhad (TNB), launched the Green Lane Pathway, which was an initiative designed to expedite

electricity supply for data centers by reducing implementation timelines from 36 to 48 months down to just 12.

But one must return to Singapore. Johor's proximity to Singapore's digital infrastructure – in particular its connectivity – and its strengths relative to Singapore make the current arrangement an attractive division of labor for both parties. In January 2025, an agreement formalizing the creation of a Johor-Singapore Special Economic Zone was signed, which provides further tax incentives for a variety of industries looking to set up shop in the area. This includes the data center industry, which seems to be broadly grouped under AI and quantum computing.

The point here is that Johor possesses a unique set of advantages over its competitors in the region. Its geography, the government's willingness to accommodate data center operators, and its ability to leverage its connection to Singapore mean that the greatest growth has remained in Malaysia and not neighboring Indonesia, Thailand, or Vietnam.

All of these countries are experiencing growth, but it seems unlikely that the center of gravity in the region is going to move anytime soon. S&P Global's Kang Soon Chen said that "customers generally like to put their resources in multiple locations. It's a concentration of risk kind of thing." In other words, one should not jump to conclusions when presented with signs of growth – this includes the recent influx of investments into Thailand and Vietnam's decision to remove a cap on foreign ownership of data centers.

Looking to the future

But all this growth does come with challenges, and Johor's main challenge – power – is hardly unique. Governments eventually have to reckon with the resentments of a population that find it unfair to use sovereign resources, be they water, energy, or foregone tax revenue, for 'development,' however nebulously defined. Data centers do not create the volume of jobs found in manufacturing or other industries, and promises of growth and development can seem empty when one's electricity and water bills start to rise.

There are signs that the government may have already overcommitted. As pointed out by a report published by ISEAS-Yusof Ishak Institute, a thinktank based in Singapore, TNB, which is Peninsular Malaysia's sole electricity distributor, owns nearly 51 percent of Malaysia's power generation market share. In 2024, it offered a total contracted capacity of 13.76GW, and up until the end of the year, a total of 38 data centers across Malaysia had secured a maximum demand of 5.9GW through Electricity Supply Agreements. In other words, 43 percent of TNB's total capacity would have been used if all facilities were to run at full throttle. Of course, actual load utilization in this period was lower than this maximum, but the growth of the market means that a greater percentage of TNB's capacity could potentially be taken up if power infrastructure is not expanded.

But, as the citizens of many townships and cities in the US have come to recognize, grid expansion of this scale is usually supported by rising energy tariffs. Malaysia is much the same. TNB announced in December 2024 that the period from July 2025 to the end of December 2027, would see the utility's allowed capital expenditure balloon to RM42.8bn (\$10bn), up 108 percent from the previous period. This will be utilized for transmission and distribution upgrades. New tariffs for that same period were also announced in December, with the tariff hike forcing data centers to pay 14 percent more for electricity.

The hike has not been met with loud disapproval. It helps that Johor's data centers are clustered around areas specifically designed for them, like Nusajaya Tech Park. Chen adds that "public sentiment against data centers, especially in Johor, is not really encroaching into the residential areas," alluding to the fact that Johor and Malaysia has yet to reckon with the demands of angry residents, however justified or unjustified. But the onslaught of data centers has yet to truly arrive, and many of the current challenges will only intensify as planned developments become a reality. That being said, as it stands, there is no large-scale opposition in Johor. It is a far cry from places like Northern Virginia or Ireland, where residents and politicians have rallied

"...a slowdown or market correction could occur if infrastructure, sustainability, and local engagement are not managed responsibly and scaled effectively,"

>> CBRE

against the creep of digital infrastructure.

Social and financial costs aside, it is also unclear whether Johor's energy infrastructure will be able to keep pace with its pipeline – 1.4GW in committed prospects depend on it. It also assuages no one's worries to know that similar issues exist with Johor's water infrastructure.

S&P's Chen noted that the government was trying to be more selective, and that several applications had actually been rejected: "The boom happened so fast," she says. "There's been a lack of planning in terms of infrastructure." CBRE, in starker terms, argues that "a slowdown or market correction could occur if infrastructure, sustainability, and local engagement are not managed responsibly and scaled effectively."

According to Johor's state government, the ad hoc Data Center Development Coordination Committee, which was formed in June of last year to vet new projects, rejected 30 percent of applications because they failed to meet certain sustainability and infrastructural standards.

These standards are also incorporated into the government's various initiatives. Since December 2024, data centers wishing to apply for tax incentives under DESAC must be 'sustainable,' which is gauged with reference to PUE, CUE, and WUE. The Corporate Renewable Energy Supply Scheme (CRESS), which was introduced in September 2024, allows new businesses to directly obtain renewable energy from independent providers through the TNB grid system.

President Anwar Ibrahim, when introducing the national budget for 2025, said that Malaysia "can no longer sustain the outdated approach of offering incentives and support to investors without considering the economic spillover.

"Investments in data centers should not be pursued unless they bring tangible added value to the rakyat [meaning ordinary people], such as high-paying job opportunities and knowledge transfer. A shift in focus is now essential, ensuring that the support provided yields economic spillovers that directly benefit the rakyat [people] and the nation, rather than merely serving the profit motives of investor companies."

Johor also faces some specific challenges. Like the rest of the region, Malaysia has needed to balance its relationship with America, the regional hegemon for the latter half of the 20th century, and China, which is the historic and rising power.

But courting interest from both countries has proven thorny in recent years. The business relationship between China and Malaysia is important – China is Malaysia's second largest export market and the country from which it imports the largest value of goods – and the presence of China-affiliated data center operators in Johor left it especially vulnerable to allegations that the country had played a significant role in the smuggling of Nvidia chips to China, or of Chinese AI models into Malaysia.

Two recent gaffes made things worse. In May 2025, the country's Ministry of Investment, Trade, and Industry had to distance itself from a speech made by the Deputy Minister of Communications,





Satellite imagery of Nusajaya Tech Park – Credit: ArcGIS Online

in which they said that Malaysia would use 3,000 of Huawei's Ascend AI chips to build the country's sovereign AI infrastructure. This would have contravened the Trump administration's AI chip export controls, which stipulate against using Huawei's Ascend chips anywhere in the world.

The following month, the *Wall Street Journal* reported that unnamed Chinese AI companies were flying suitcases of hard drives to Malaysia in order to use Nvidia chips to train their AI models. This also goes against American export rules.

The month after, fear of total cutoff prompted the country to introduce a permit system for the movement of US AI chips. So far, the country seems to be in the clear. But even if the worst-case scenario does occur, many of

those who have secured contracts to build data centers already possess chips, meaning the impact would be most acute on Johor's future prospects rather than those already locked into the development timelines.

There is also the broader question of whether an AI bubble exists, and how Johor might fare if it pops. The answer is unclear. According to a survey conducted by Microsoft and LinkedIn about the state of AI at work, 84 percent of Malaysian knowledge workers had used AI at some point in their workflow. But it was not clear whether this statistic encompasses those workers who try ChatGPT out once, or whether it refers to individuals who consistently use AI.

CBRE has said that "demand is driven by the real need for AI and cloud workloads, and strategic positioning

after the Singapore moratorium has also fuelled regional AI infrastructure growth", but S&P's Chen adds that when it comes to AI adoption in the public and private sectors, it is still "an early stage for a lot of developing economies."

Much of this comes back to the dynamic that defines data center buildout across the world. Because of various geographical, financial, and political reasons, the place where the computer is built is no longer the place where it is being used. China has 'Eastern Data, Western Computing,' the US has its water-stressed Midwestern rural towns; Europe has Ireland; and now Singapore, and the rest of the world, has Johor. And if domestic AI adoption turns out to be meek, residents may well start to question whether all the infrastructure being built in Malaysia actually benefits Malaysians. ■



Grundfos Data Center Solutions

Efficient cooling that makes an impact

Optimised water solutions
for efficient data flow

What if you could maximise your data center's performance while minimising the environmental impact?

Innovative Grundfos Data Center solutions make it possible. From air and liquid cooling systems to water treatment and reuse, our solutions and services give you the best chance to achieve your sustainability targets. With a 75-year history of sustainability, quality and innovation, you can rely on Grundfos to keep your servers cool, allowing you to enjoy a hassle-free workday with no worries.

Experience Grundfos Data Center solutions at
grundfos.com/datacenters

GRUNDFOS 

Possibility in every drop

The big fightback



Paul Lipscombe
Telecoms Editor

Regular attacks against telecoms infrastructure have got the networks fed up

With smartphones dictating every facet of our lives, connectivity has never been more crucial.

The sending of messages, making payments, and surfing the Internet mean telecoms infrastructure has fast become critical infrastructure. So, it's almost unfathomable that some would attack this vital equipment.


However, this is becoming a reality for a lot of telcos and communications providers at present. And simply put, the telecoms industry is fed up with it.

The problem has become so bad that some

within the industry have called the issue an "epidemic" and a threat to national security.

"The reality is that there is no shortage of individuals, rogue groups, nation states, and others seeking to cause harm to Internet users, broadband networks, and in some cases, the collective critical communications infrastructure," said former US Federal Communications Commissioner (FCC) Michael O'Rielly earlier this year.

As he says, these attacks can be for several reasons, such as people deliberately vandalizing telecoms infrastructure, stealing critical



equipment such as copper cables, or even sabotaging fiber cables.

Copper cables in particular draw a large amount of interest, given the street value of copper.

Rising numbers

In the past year alone, vandals have caused millions of dollars' worth of damage to telecoms infrastructure.

According to Protect Critical Communications Infrastructure (CCI), a coalition designed to "protect America's critical communications infrastructure from theft, vandalism, and the costly disruptions they cause," in the last six months of last year, nearly 6,000 intentional incidents of vandalism against communications infrastructure were reported.

In one instance, suspected copper thieves toppled a 488-foot rod cell tower in Oklahoma, causing an estimated \$500,000 in damage, while just this summer, a man was arrested after stealing more than 900 backup batteries used at cell towers across Phoenix, Arizona. The latter targeted 66 T-Mobile cell tower sites between August 2024 and February 2025, resulting in more than \$330,000 in damages.

Bell's copper conundrum

Bell Canada is one carrier experiencing more incidents. Copper theft-related incidents have increased 23 percent year-over-year across Canada as a whole, and since 2022 Bell alone has recorded more than 2,270 thefts across the vast country.

"Copper theft is a crisis that continues to grow," David Joice, director of network operations, Bell Canada, tells *DCD*. "It endangers public safety and disrupts critical communications in Canada as well."

Joice explains that the issue is showing no sign of improving, with numbers rising even higher, stating that the problem is costing the company "millions of dollars."

So far this year, Bell has reported 700 such incidents, with around 88 percent down to copper theft. June saw 200 incidents of vandalism against its network. "It just continues to escalate," he adds.

The key to this issue is the black market that is driving demand for copper, explains Joice.

"When the impact of these intentional attacks impacts public safety, it isn't just vandalism - this is a form of domestic terrorism,"

>> Charter statement

"The cost of metals has increased over the years, and people are not only stealing it, but people are willing to buy it, knowing that it could be stolen," he says.

The situation is similar in the US says Chris Antlitz, principal analyst, Technology Business Research. He likens the issue to one that was prominent for Toyota Prius car owners, whereby thieves targeted the vehicles for their catalytic converters, due to the high scrap value of the materials.

"It's a big problem in the US. The cost of materials has reached a point where illegal activities get a payout that's worth their time," he says.

"People used to take the catalytic converters from the car to reclaim the metals. It's the same idea with the telecom infrastructure. There's a lot of copper in telecom equipment, especially legacy stuff."

Joice acknowledges that attacks are an everyday battle for telcos such as Bell.

"I'd say we don't catch the majority of them. We do catch some and I'd say over the years, it seems like we've gotten better at catching people, but the fact that it continues to happen every day, it seems like we're not winning that battle," says Joice.

He notes that Bell works with law enforcement on almost every case to prosecute criminals, but says the vast landmass of Canada can often mean it's difficult to keep on top of incidents, noting that attacks are often in rural areas.

"An act of domestic terrorism"

US-based Charter Communications has also been the target of vandals.

However, unlike Bell and other carriers, Charter's network doesn't include copper, but is made of fiber. This has not deterred criminals, despite fiber not having the value of its metal counterpart.

An attack on Charter's fiber network in Los Angeles in June led to an outage that impacted a range of customers, including a US military base, emergency dispatch and 911 communication services, local fire and police departments, financial institutions, carriers, and more. On that occasion, 13 fiber cables were severed, including more than 2,600 individual fibers.

The company labeled that incident as an "act of domestic terrorism." Sadly, it wasn't the only such incident for the company, with Charter outlining a recurring issue in Missouri, another state, close to 2,000 miles away.

Up to the end of June, Charter logged 148 outages related to criminal attacks on infrastructure in the state of Missouri alone.

"The impact is that of domestic terrorism, cutting off access to critical services, preventing vital communications during times of emergency and crisis, creating a pervasive and persistent threat to families and businesses in Missouri and across the country. This criminal activity must not be allowed to continue," said Tom Monaghan, executive vice president, field operations for Charter in July.

In a statement to *DCD*, the company has doubled down on its view of these attacks against telecoms infrastructure.

"When the impact of these intentional attacks impacts public safety, it isn't just vandalism - this is a form of domestic terrorism," says a Charter spokesperson.


They added that stealing fiber does confer the same financial benefit as copper.

"A driver of this issue is the rising price of copper and precious metals," the spokesperson says. "Many of these attacks are being perpetrated to steal copper often found in cables. However, our fiber networks contain no copper, and so criminals are cutting these lines, not finding any copper, but causing significant damage and impacting the services and in many cases, public safety."

Misinformation

The attacks against the infrastructure are not always in the pursuit of a quick payday. In some instances, people have attacked infrastructure due to the spread of misinformation online.

One of the biggest examples of this



was around the pandemic back in 2020, when numerous attacks across parts of the world saw people destroy or attack 5G cell towers as part of "Anti-5G" protests.

This was down to conspiracy theories being shared online incorrectly linking the spread of the virus with 5G technology, and even sometimes pedalled by celebrities.

Incredibly, five years on, several 5G masts were targeted in parts of Belfast, Northern Ireland earlier this year.

Protests around mobile technology are nothing new, says Ineke Botter, who has held high-level positions across several countries in Europe, plus CEO positions at telcos in Kosovo, Lebanon, Azerbaijan, and Haiti.

In one instance, Botter recalls people chaining themselves to 2G cell masts in Switzerland amid fears of health concerns, though these protests eventually petered out.

She says the misinformation can be "dangerous" and has to be nipped in the bud quickly.

"You have to be vigilant of social media," she says. "In Haiti, you saw people stirring up nonsense on social media, and you've seen the same with things like '5G is dangerous' and all that bullshit," explains Botter. "If you don't manage that constantly and you don't nip it in the bud you have a real problem."

That said, her role as a CEO in war-torn countries presented trickier assignments. In Lebanon, the terrorist group Hezbollah was a regular threat to the communications infrastructure in the country.

"Hezbollah had a sort of attitude that

"We are already coming off the copper network nationally, across the UK, and we are extracting copper at a fast pace, almost in a race with the criminals, because the quicker we can get copper out of the ground, the quicker they can't get hold of it,"

*>> James Arnold,
Openreach*

everything is ours," explains Botter. "They were stealing infrastructure, such as cables, and were stealing from the generators. So what you had to do there was to make sure that the plots were very well protected with dogs and guards and all, but of course, you don't win against Hezbollah."

Attacks slow down fiber build-out

UK broadband provider Openreach has experienced its own fair share of issues with vandalism against its network, in particular, copper theft.

James Arnold, regional security operations manager, UK South & London at Openreach, told DCD that theft of its copper network has impacted the business, especially as the company pushes to build out its fiber network. The company plans to reach 25 million premises with its fiber network by the end of next year.

"Cable thefts are hugely disruptive. The loss of phone and broadband is not only inconvenient but can put vulnerable people at risk. Repair work also pulls our engineers away from other work, can take weeks to finish, and costs thousands of pounds," says Arnold.

"I'm dealing with a welfare case whereby a vulnerable customer has been impacted by cable theft in Lincolnshire," he adds. "It can detract from important work, such as our fiber rollout. It detracts our engineers because they then have to do the repair."

As copper theft has become such a big issue for Openreach, it's led the broadband provider to step up its security to prevent copper from being stolen from its network, which can lead to network outages, such as in the New Forest, UK, last year.

To combat such theft, Openreach has used tracking technology, SelectaDNA, to trace its copper, noting that the technology has seen copper theft drop by 30 percent during the last year.

"We take the security of our network seriously and have a wide range of crime prevention tools to prevent thefts and catch those responsible. Our dedicated security team investigates all attacks, and our network is alarmed and monitored 24/7 by our control center," he adds.

It does, however, seem to be working, says Arnold, who notes Openreach has seen a reduction of 37 percent in live network attacks year-on-year, while the company has reported a 90 percent decrease in estate and cable compound targeting.

Unfortunately, it's not just the infrastructure that has been attacked; but Openreach has also reported a significant amount of abuse directed towards its engineers.

Last year, the company logged 450 incidents of abuse and assault on its workers in the year up to the end of March, up eight percent year-on-year, and up 40 percent up from 2022-23.

"We've had engineers being shouted at, sworn at, even spat at," says Adam Elsworth, health and safety director at Openreach. "We've had engineers where their vans are blocked on people's drive and even people being shaken off of step ladders or pushed down people's stairs."

He notes that around half of the incidents occur in the public domain.



"You need to really make people understand that these base stations are there for them, so that there is communication. If you destroy those, you're out of communication,"

*>> Ineke Botter,
former telecoms CEO*

In an effort to support their workforce in such incidents, Openreach is utilizing a "panic alarm" on engineers' mobile phones, which connects them in seconds with a monitoring center. The center has the ability to directly dispatch emergency services if needed.

Tougher punishments are needed

The issue at hand is a big one for telcos globally, and it needs to be addressed.

Telecoms associations are becoming more vocal in calling for tougher laws to curb these attacks.

In July, four US telecom trade associations wrote a letter asking for urgent action to be taken to safeguard the nation's critical network infrastructure.

In the letter, which is addressed to Kristi Noem, secretary of the US Department of Homeland Security (DHS), and Kash Patel, director of the FBI, the Internet & Television Association, ACA Connects, NTCA – The Rural Broadband Association, and USTelecom, calls for tougher measures and action to prevent further attacks.

"We write today to highlight a significant and rapidly growing threat demanding urgent, coordinated federal, state, and local action: widespread, organized critical infrastructure attacks targeting our nation's communications networks," said the four associations.

DCD tried to contact the associations for more information on several occasions, but got no response.

Stepping up security

The telecoms companies have taken it upon themselves to combat the issue at hand by bolstering security.

South African telcos, including MTN

Group, Vodacom, and Telkom, have previously strengthened their security to deter vandals and thieves, while also working closely with police forces to clamp down on the issue.

Joice says Bell is doing the same in Canada, noting that security has become a key priority.

"In some cases, we actually have alarms on the copper cables. So if they are cut, that notifies us and the local authorities right away," says Joice, who adds that this has enabled the police to apprehend people faster.

"We have lots of copper out there, so we have to be fairly strategic where we place these systems. We've also installed more security cameras and hired security guards."

The company hopes its move away from legacy copper networks will lessen the impact of the vandalism.

Meanwhile, Charter is also working closely with the relevant authorities to detain vandals.

"We are able to quickly identify when and where these attacks occur and quickly dispatch our teams to begin repairs," says Charter. "To fight this issue, we are working closely with law enforcement agencies and prosecutors to identify, catch, and prosecute those responsible."

The company says it has also worked with elected officials to pass legislation seeking to increase the penalties for these crimes. "In 28 states these critical infrastructure attacks are now a felony offense," its spokesperson adds.

Take control of the network

While the need for tougher sentences is widely agreed upon by the industry, some suggest that there's a need to educate the

public on why this infrastructure is so crucial.

According to Antlitz, there's an opportunity for telcos to take control of the copper issue, by reclaiming their copper before opportunistic vandals can.

"The telcos are reclaiming their own copper," he says. "They are ripping out their own copper anyway to recycle it, because a lot of that equipment is legacy, and they're replacing it with fiber."

"The fiber doesn't have value on the secondary market, so they're not really stealing the fiber, unless there's vandalism."

Arnold agrees, saying the company wants to get its fiber in the ground as soon as possible.

"We are already coming off the copper network nationally, across the UK, and we are extracting copper at a fast pace, almost in a race with the criminals, because the quicker we can get copper out of the ground, the quicker they can't get hold of it," he says.

Botter believes that educating the public would potentially deter more people from attacking the infrastructure that plays such a key role in everyday life..

"If you can't explain difficult things in very simple terms, then you are not good for your job," Botter says.

"You need to really make people understand that these base stations are there for them, so that there is communication. If you destroy those, you're out of communication."

As she points out, when war or natural disasters strike, the power of telecoms infrastructure becomes more critical and non-negotiable for everyone, which should give wannabe telco criminals food for thought. ■

CORNING

GLASSWORKS AI™

**Leverage our AI expertise
to build yours!**

GlassWorks AI™ reimagines the plan-design-deploy process for you to create your customized AI network.



Scan the QR Code
to Unlock New AI Possibilities

© 2025 Corning Optical Communications.
All rights reserved. LAN-3472-A4-AEN / July 2025

HITACHI

See **what's next** in
clean energy for
data centers

Hitachi Energy provides
stable, clean power to the
world's data centers.



Explore more

Inspire the next



How to build a supercomputer



Charlotte Trueman
Compute, Storage,
and Networking Editor

DCD gets up close and personal with Hunter, the most powerful supercomputer at the Stuttgart HPC center



When you hear the term ‘supercomputer,’ for the uninitiated, it’s likely to conjure images of a machine akin to the Colossus computer used by British code breakers at Bletchley Park during the Second World War.

And, while some modern-day systems are undoubtedly impressive in their size – Elon Musk’s efforts to develop his own Colossus cluster have thus far seen him deploy more than 200,000 GPUs in Tennessee – for most systems, visually at least, the reality is far more underwhelming.

Despite the best efforts of British Airways to keep us grounded, DCD traveled to the High-Performance Computing Center (HLRS) in Stuttgart, Germany, to visit Hunter, the first of two HPE supercomputers ordered by the University of Stuttgart to take HLRS up to exascale level.

Brought online in January 2025, the €15 million (\$17.5m) HPE Cray EX4000-based system comprises 512 AMD Epyc ‘Genoa’ processors, with the CPUs grouped into 256 nodes, each equipped with 768GB of DDR5-4800 memory. However, unlike other HPC systems housed at HLRS, Hunter is also powered by GPUs: 752 AMD Instinct MI300A accelerated processing units (APUs) across 188 liquid-cooled nodes.

Officially launched in December 2023, AMD’s Instinct MI300A combines 24 Zen4-based CPU cores with a GPU accelerator and 128GB of high-bandwidth memory (HBM3) on the same silicon. This architecture allows both the CPU and GPU to access the same memory, which AMD says saves trillions of calculations from going through unified memory and speeds up the performance of HPC workloads.

Each APU has 24 Zen 4 CPU cores and offers 128GB of HBM3. AMD has also used a 3D stacking design for the chip, allowing it to integrate several smaller silicon dies to form a larger processor.



Hunter has a theoretical peak performance of 48.1 petaflops (more on that later), and each of its nodes is equipped with four HPE Slingshot high-performance interconnects. It ranked 54 on the most recent edition of the Top500 list, and 12 on the Green500, and provides double the speed while consuming 80 percent less energy than its predecessor, Hawk.

When it comes to Hunter’s architecture, Utz-Uwe Haus, head of HPC/AI EMEA research lab, at HPE, describes the Cray EX design as “the architecture that HPE, with its great heritage, builds for the top systems.”

A single cabinet in an EX4000 system can hold up to 64 compute blades – high-density modular servers that share power, cooling, and network resources – within eight compute chassis, all of which are cooled by direct-attached liquid-cooled cold plates supported by a cooling distribution unit (CDU).

“It’s super integrated,” he says. “The back part, which is the whole network infrastructure (HPE Slingshot), matches the front part, which contains the blades.”

For Hunter, HLRS has selected AMD

hardware, but Haus explains that with Cray EX systems, customers can, more or less, select their processing unit of choice from whichever vendor they want, and the compute infrastructure can be slotted into the system without the need to total reconfiguration.

“Should HLRS decide at some point to swap [Hunter’s] AMD plates for the next generation, or use another competitor’s, the rest of the system stays the same. They could have also decided not to use our network – keep the plates and put a different network in, if we have that in the form factor. [HPE Cray EX architecture] is really tightly matched, but at the same time, it’s flexible,” he says.

Hunter itself is intended as a transitional system to the Herder exascale supercomputer, which is due to go online in 2027. A new data center is currently under construction at HLRS ahead of Herder’s planned deployment, as reinforcing the floor in the data hall where Hunter is housed to have the two systems sit alongside each other is not possible.

Despite all this, the colorful and rather powerful system looks a little lost sitting in the corner of a data hall, surrounded by all the empty space that Hawk used to occupy, and the center’s Vulcan NEC cluster – a “standard PC cluster” available to the University of Stuttgart for research, visualization, AI, and big data workloads.

Simply the best

HLRS was established in 1996 as Germany’s first national high-performance computing (HPC) center, providing users across science, industry, and the public sector with access to supercomputing resources for “computationally intensive simulation projects.”

It is one of three national computing centers in Germany, the others being the Leibniz Supercomputing Centre (LRZ) in Garching, near Munich, and the Jülich Supercomputing Centre (JSC) in Jülich, North Rhine-Westphalia. There are also a number of smaller supercomputing centers in the country, which are overseen by the Gauss Alliance, a non-profit association for the promotion of science and research, and for which Prof. Dr. Michael M. Resch, director of HLRS, is a board member.

Getting a supercomputer like Hunter

“Peak performance, ladies and gentlemen, is like taking a car, putting it on a plane, and then taking off and throwing the car out of the window and then measuring the speed of the car. That is a ridiculous number,”

>> Prof. Dr. Michael M. Resch, Director of HLRS

from concept to reality is a process that requires both planning and patience, in addition to expertise, a long-term business model, and an investment structure that includes a whole ecosystem of hardware, software, and infrastructure to go alongside it.

Resch says HLRS has been deploying systems comprising AMD hardware since 2002, yet, every time the center starts to discuss building its next system, a vendor-agnostic approach is taken – when asked if Herder will also be comprised of AMD hardware, Resch says no decision has been made yet, and negotiations will be ongoing until the end of 2025.

"We are a national supercomputing center and we need to buy a new system every now and then," he says, adding that when this happens, the center goes through a rigorous procurement process, which involves lots of long discussions with vendors.

"We went for the best solution. I was asked... whether European was important, and the answer is no. We want to have the best system, and we don't care [where it comes from] – we are very agnostic when it comes to company names or nationality. We have around 800 users out there, and they need the best solutions."

Resch says that in the world of HPC, the best solution means the one that offers sustained performance.

"From time to time, I get the question 'Why is your peak performance... not as high as this or that [system]?' And I say: 'I don't care.'"

"Peak performance, ladies and gentlemen, is like taking a car, putting it on a plane, and then taking off and throwing the car out of the window, and then measuring the speed of the car. That is a ridiculous number, much as peak performance is ridiculous, and the colleagues here, they know that, but not everyone in the market knows that. Peak performance is not relevant; the question is 'how much do you get of this performance?'"

Consequently, HLRS has eschewed traditional benchmarks and instead sets vendors a test that consists of running three of the center's production codes.

"This is much more exciting," Resch jokes.

Getting a supercomputer like Hunter from concept to reality is a process that requires both planning and patience, in addition to expertise and a whole ecosystem of hardware, software, and infrastructure to go alongside it.

In addition to sustained performance, cost is a big factor. As Resch says, it's all very well and good for the President of the European Commission to say there will be hundreds of billions of euros worth of investment in a certain area when ultimately, they aren't putting up the money or paying the final bill.

He explains that there's always a push and pull between "make or buy" when it comes to weighing up these costs.

"When you have your own supercomputer, you have a huge investment. You're spending tens of millions just to make sure that the system is up and running, and then, after five years, you spend another 10 million to buy a new one.

"But, on the other hand, the advantage is that you have a higher level of flexibility, in the sense that it's your own decision of what happens with it. [At HLRS], we will buy a system every five years or so, [but if users] want to have another system in two years, we will say, 'Sorry, we don't have another system. You'll have to wait for another three years.' So, that level of flexibility is there."

However, when it comes to a question of cost and profit, because HLRS is a national center, it can only charge its users what it costs for the center to run the system – "no more, no less" – as doing anything different would be illegal. HLRS is also not allowed to provide subsidized compute to any particular user or institution.

Consequently, HLRS is cheaper than cloud providers, particularly when it comes to offering compute to non-cloud native workloads that are tightly coupled or require longer durations, as it does

not have to take expenditure recovery or profit margins into account when setting its pricing strategy.

Furthermore, while commercial hyperscalers must provide availability guarantees of 99.99 percent as part of their service level agreements, HLRS does not have such an obligation, meaning the center does have to spend money on UPS (uninterruptible power supply) units or generators.

"Maybe there are hospitals using [an unspecified cloud provider], so it cannot have a power outage. [The provider] must therefore have an uninterruptible power supply, have diesel engines, and have all kinds of things to make sure that it is absolutely safe.

"We don't do that, and we can't afford to do that because it's public money that no government will accept. If I said: 'We are at 99.8 percent, but for the extra 0.2 percent, I need an additional €50 million (\$58m) for the power supply, for the diesel,' the center would say 'no, we don't need that.' That's something which our customers accept."

For this reason, Resch says HLRS will not accept requests from institutions such as hospitals, banks, or insurance companies as they cannot guarantee the level of uptime needed by such customers.

The future of HPC

While Resch said that no final decisions had been made yet regarding Herder's likely hardware, if the center were to select AMD again, the supercomputer would be in highly esteemed company.

On the most recent edition of the Top500 list, around 34 percent of supercomputers were AMD-powered, with the chip company also responsible for powering the top two exascale systems on the list, El Capitan and Frontier, which took the first and second spots, respectively.

The June 2025 edition of the Top500 also represents the seventh consecutive list on which the world's most powerful supercomputer has been powered by AMD.

However, despite his evasiveness about Herder, during the visit, Resch did let slip the existence of a previously unannounced AMD MI600 AI chip, demonstrating just how advanced the planning is for AMD's hardware

offerings, and that the company is already providing customers with technical details for chips that are multiple years out from production.

"In 2023, we signed a contract for the year 2027, and that was a risk on the vendor side and on our side."

Resch says that as part of that agreement, AMD and HLRS agreed they would sit down in the last quarter of 2025 to discuss the path going forward, and have the chip maker disclose the specifications of the chip it's planning to bring to market in 2027.

"Four years ahead of delivery, we don't get detailed specifications, so now, we need to look into these detailed specifications, talk to HPE about it, and find out if this changes anything with regard to the overall system. I believe these negotiations will be very short and very nice, but right now, I cannot say in advance what the best solution is."

Peter Rutten, global research lead on performance-intensive computing solutions at IDC, says that this is one area in which AMD has been "incredibly good" – the execution of its roadmap.

He adds: "With every new deliverable stage, AMD has basically outdone the expectations of its new generation processor, and so the market has just come to expect that what AMD does is always going to be spectacular, and they so far haven't failed on that."

Rutten says that while it's not surprising AMD has continued to see success in the HPC space, it's still a relatively new phenomenon for the company, with its traction in the market only starting around five to seven years ago.

There are a number of factors that have – and continue – to contribute to AMD's dominance, one of which is cost, attributed to the company's ability to provide high performance and efficiency, which is crucial for HPC centers with limited budgets.

"With HPC, the rule is the best performance in the most efficient way," Rutten says. "Think of an HPC center. Think of their budgets. Think of what they're trying to do. They're trying to solve very tough scientific problems, but often, they have limited budgets, usually government or academia. So, how do you get the most performance

with the least amount of expenditure? That is always the critical question for any HPC site, and AMD responded to that question convincingly better than their competition."

One of the ways in which AMD was able to achieve this was by being the first company to manufacture 7nm technology, which they combined with a redesign of their processors in a way that made them more performant. Rutten says that as a result, the company was able to become a market leader in the HPC space, particularly with regard to intensive and technical engineering workloads, whilst also allowing it to establish itself as a player when it came to AI.

That's not to say other companies aren't also seeing success in the HPC space. GPU giant Nvidia, a company that has become synonymous with AI, is currently powering four supercomputers



in the top ten of the Top500, with 70 systems listed deploying the company's H100 GPUs.

Rutten also notes that Intel, which has had a turbulent couple of years from both a financial and hardware perspective, is still very much the market incumbent, powering the third-place exascale system Aurora and providing CPUs for the largest share, 294 systems or around 59 percent, of Top500, though this number is falling.

For AMD, that figure sits at 173 systems, around 34 percent, a modest increase from 32 percent six months ago.

Additionally, a total of 237 systems on the list are using accelerator or co-processor technology, up from 210 six months ago, an architecture that Rutten says has played a significant role in Nvidia's success, particularly as an increasing number of AI supercomputers are deployed.

AI processing comes from the GPU's ability to process things in parallel, as opposed to with HPC workloads, which often just require very fast serial processing, but not necessarily parallel. This meant that traditionally, server nodes didn't come with a co-processor as standard, something Rutten says has now changed because it's been shown that HPC workloads really benefit from that parallelization that GPUs provide.

"A server with a GPU in it also needs a processor, you can't run a GPU without a processor also being present," he explains. "Intel was having a decent market with servers that ran on Intel CPUs and Nvidia co-processors, but that has been changing because Nvidia has seen that AMD is also developing very performant processors, which is why in recent years, we've seen more and more servers with AMD processors as well."

"Nvidia has been very focused on AI, but, along the way, realized that HPC was also a very attractive workload for them," Rutten says. "Initially, that was sort of an afterthought, but then it became actually an adjacency for them to focus on, and they have – to the point where a lot of supercomputers are now Nvidia GPU accelerated."

While Rutten argues that Nvidia isn't doing anything totally revolutionary, rather just taking a what is "essentially a newer approach to HPC," he says that the old way, with just CPUs and tightly connected server nodes, had a lot of practitioners that were very skilled in optimizing HPC environments, one of the most difficult but critical jobs in any supercomputing lab.

"There are people who have PhDs, but with these GPUs now becoming part of the supercomputers, that skill set has changed, and that traditional way of thinking is being challenged. There are now different considerations going into how you optimize a supercomputer, rather than just what was involved with a non-accelerated supercomputer."

"I don't think we have fully gotten to the point yet where the skills to optimize an accelerated supercomputer are as advanced as the skills to optimize a classical supercomputer, although we're getting there. But, it has been a little bit of a challenge for HPC sites to understand what they could do with GPUs and how to do it." ■



／ Engineered For Impact

What we offer

As your end-to-end partner, Salas O'Brien delivers high-performance mission-critical projects through pre-engineered solutions and offsite prefabrication that reduce risk, limit rework, and accelerate results.

- ▲ Pre-engineering & offsite prefabrication
- ▲ Mechanical, electrical, plumbing, and fire protection
- ▲ Architecture & interior design
- ▲ Technology & telecom systems
- ▲ Decarbonization & sustainability
- ▲ Commissioning & QA/QC

**Let's talk about your
next project.**



COMPREHENSIVE ENGINEERING EXCELLENCE



www.anordmardix.com



A VFX studio's story, Untold



Georgia Butler
Cloud & Hybrid
Editor

How Untold Studios grew from the ground up, fully on the cloud

In 2018, Amazon Web Services was given a proposal. A visual effects (VFX) studio wanted to launch, and it wanted to do it entirely based on the cloud, a feat that, up until then, had never been done before.

Untold Studios was founded by Rochelle Palmer, Darren O'Kelly, and Neil Davies with the key philosophy of hiring the most "exceptional creative people" regardless of where they were located, and that meant not being tied to any particular physical location - or data center.

"Initially, when we spoke to AWS about

it, they said it's never been done before to do that kind of high fidelity imagery in the cloud, and it just wasn't possible," Palmer says. "We said, 'we want to do it,' so they built a lab for us."

"We spent around three months testing at the beginning, and then we opened. There were lots of people who didn't believe we were fully cloud, they would come and look for secret rooms in the back," she adds, laughing.

Today, Untold is an established studio with an impressive array of Hollywood projects under its belt, with VFX work for blockbuster movies including *Mission:*

Impossible - Dead Reckoning and *Sonic the Hedgehog 3* showcased on its website, alongside work for many other clients.

Amaan Akram, CTO of Untold Studios, tells *DCD* that part of the draw of the cloud is that the company doesn't want to be bogged down by its technology. While it has physical spaces and "studios" for staff in London, LA, and Mumbai, those facilities focus on providing a "physical space for employees" because "people need to see people, and we want to make sure that our creatives are sitting together and working together as much as possible."





VFX and the cloud

The concept of rendering and creating VFX in the cloud is growing in popularity. DCD has reported on the topic in the years since Untold launched - with Wētā FX using AWS to render the compute-intensive movie *Avatar: The Way of Water*, after an entire data center was not enough to handle the workload.

Cloud providers have added media and VFX-targeted offerings. AWS has its "Deadline Cloud," and in April 2024, Akamai launched a cloud-based offering specifically for video processing.

That solution was based on Nvidia RTX 4000 Ada Generation GPUs, but John Bradshaw, Akamai's field CTO, explains that while the GPUs are good for AI workloads, the company also offers a "VPU," video processing unit, which is much better for the job.

"These are custom-built pieces of silicon that are designed to process video at massive rates, whether you are transcoding or changing the format, whatever you are doing - you can put a lot more through these cards than you can a GPU, and for a fraction of the price," Bradshaw says. "When you get to a large scale, that starts to have an impact because if you are a VFX studio, you are burning through an awful lot of cycles."

Akram says Untold Studios mostly

"The intersection of generative AI and CGI image creation is the next wave. The convergence of these two totally different engines of creation will unlock new ways for artists to tell stories in ways we could have never imagined,"

>>>James Cameron,

Director

relies on "many, many CPUs" on AWS, though GPUs are used when the company turns to AI.

"There are instances where we use both, and in the future, I think it will probably become a more equitable basis for CPU vs GPU," he adds.

It uses AMD and Intel chips - as the company's software isn't built for one specific CPU configuration - accessed via

AWS G4 and G6 instances, among others. According to AWS, the G instances are ideal for "graphics-intensive applications such as remote graphics workstations, game streaming, and graphics rendering," and can also work well for AI workloads.

Notably, these instances do not feature Nvidia all-conquering AI chips. Akamai's Bradshaw says this is probably because "AI-driven processors have ginormous back planes, so their ability to move gigabytes, or terabytes, of data in and out in fractions of a second is really important, because if you have an AI workload, you have to load the vector database in, and that contains all the relationships and the calculations you would run upon it. That is not very useful for video processing.

"By and large, GPUs are great to bring data in and out very quickly, and have huge capacity, but you're not doing the same kind of maths in these cards as you do when you're doing graphics or visualization."

The company uses Nvidia GPUs for some of its VFX work, and the H200s for some AI work - adding that in the future it expects to use AWS' own-developed chips for "internal AI needs."

Beyond the G instances, Untold will tap into "spot instances" that can include anything from the C5, C6, M5, M6, R5, R6,

or Z1 based on the needs of the project.

AI-generated content

The use of AI in the world of VFX and “high-end entertainment media” is becoming increasingly prevalent, though not without its challenges.

Stability.AI is known for its “Stable Diffusion” text-to-image tool, but has since progressed to AI-created video. Notably, in September 2024, James Cameron - the director behind VFX-fueled movies including *Avatar* and its sequel, among many others - joined the board of directors at Stability.AI.

At the time, Cameron said: “I’ve spent my career seeking out emerging technologies that push the very boundaries of what’s possible, all in the service of telling incredible stories. I was at the forefront of CGI over three decades ago, and I’ve stayed on the cutting edge since. Now, the intersection of generative AI and CGI image creation is the next wave. The convergence of these two totally different engines of creation will unlock new ways for artists to tell stories in ways we could have never imagined.”


That same month, Lionsgate - the studio behind the likes of *John Wick* and *The Hunger Games* - entered into a partnership with Runway, giving the AI company access to Lionsgate’s archive to develop an AI model exclusively for the Hollywood studio. Runway has previously partnered with AWS to train its models.

Another AI startup targeting video generation is Moonvalley, which was co-founded by former DeepMind researchers who worked on Google’s own video-generation model, and signed a multi-year contract with CoreWeave to use Nvidia GB200 NVL72 systems to train its models earlier this year.

But the use of AI-generated content comes with important considerations.

Issues remain around the ownership of the content. With AI models trained on existing creative works - that are human-made - copyright laws become a complex topic to navigate, and certainly one that Hollywood does not appear to want to deal with.

AI startups offering these services often make guarantees that all content used for generation is fully licensed, but this is not always enough for media executives.




David Colantuoni

“The people we are working with want absolute control over image quality; they will obsess over pixels

*>>Amaan Akram,
Untold Studios*

“People create the artwork that goes into these models,” Untold’s Akram says. “On top of that, clients do not want us using data that comes from foundations they are not sure about. We’re very sensitive to their requirements as well, and we want to make sure that we don’t upload work to any of these many generative AI companies, as we don’t know how it will be used, and it might end up on some other customer’s generated content.”

As a result, Untold has its own models developed from its own content. This is another area where Untold benefits from its close relationship with AWS. Both Palmer and Akram describe themselves as “creatives” first and foremost, and that training new models is the domain of AWS. “We need the real experts working on that process,” Akram says.



Untold’s use of AI seems to be more so in the minutiae of VFX, rather than replacing the actual artists. It is currently working on a “de-aging model” to make actors appear younger.

“The people we are working with want absolute control over image quality; they will obsess over pixels,” says Akram.

Palmer adds: “We’re talking about milliliters of fluid in an animal’s eye or something. It’s that sort of detail, and we aren’t in a space where you get given an image and it can’t be changed.”

This level of detail makes creating convincing VFX very difficult. Palmer says one of the most common complaints from artists is about virtual water. “It’s very difficult to CGI and not have it look like jelly,” she says.

She recalls a particularly challenging project making adverts for Virgin Media: “They wanted creatures doing things they wouldn’t ordinarily do,” Palmer says. “The first thing we did was a highland cow on a motorcycle, and then a goat on a hang glider. Virgin asked; ‘What would be the most complicated thing to do?’ And we said probably a walrus - with the blubber and water and how it all moves - and we did that on a speed boat.”

Storage and scalability

Compute complexity is also challenging for VFX studios. Untold worked on a series of adverts for the Super Bowl, an effort that took a massive amount of processing power.

“It was five huge commercials with effects that all needed a lot of complex computing power, and had to be delivered on the same day,” Palmer says.

Untold usually works on many projects simultaneously, but these are typically at different points of development and do not need the same amount of compute. This was not the case for the Super Bowl. The studio calculated that to render all five of the commercials on one artist workstation would have taken 138.4 years. Instead, via AWS, Untold was able to leverage as many as 1,470 EC2 nodes to get the work done in time.

This scalability, often touted as a key benefit of cloud computing, is particularly useful in VFX, and applies not only to the compute needed for rendering and other workloads, but also for storage.

Untold has “petabytes and petabytes,” and while Akram says the company is careful with what needs to be kept “live,” it is still a huge trove of content to manage.

Even within the cloud, physical location is an important consideration. Untold operates in multiple cloud regions that are logically selected for proximity to its studios.

Akamai’s Bradshaw says location is also important from a provider perspective. The firm’s origins as a CDN vendor have given it a disparate footprint of 4,200 locations, and it would be wildly impractical to deploy its VFX and media cloud offering at each location, not least because of how expensive the required hardware is.

“We focused on our core data centers and cloud locations,” Bradshaw says. “I think there are five or six in Europe, and we chose the locations as they also contain our object storage offering, which makes it much easier to pull the file out and do whatever you need to do.”

While Akamai caters for a variety of media needs, including broadcasting and streaming, Bradshaw says location is particularly important for VFX.

“You get these rushes and you need to process things very quickly so that a producer or director can make their judgment call,” he says. “These are ginormous files - they could be 4K, 8K, or more video, and you need to be able to access them quickly. If that needs to travel halfway around the world before it can be edited, it’s going to be an uncomfortable experience.”

While Untold has embraced multiple locations, other studios take a different approach.


Shadow Magic Studios is a cloud-based production company with an integrated digital lab services offering.

The studio isn’t handling VFX workloads, but Jordan Maltby, its founder, says it faces similar challenges. Maltby says the company’s decision to hit the cloud was motivated by a desire to bring people and talent from any location in the world.

“The first step was ‘How do we get to the cloud?’” he says. “Then, you start realizing that the cloud is in data centers in specific regions. Even with top-tier cloud providers, if it’s stored in New York and you do a pick-up shoot in California,

“You start realizing that the cloud isn’t just in the ‘ether of everywhere,’ it’s still in data centers in specific regions,”

*>>Jordan Maltby,
Shadow Magic Studios*



Jordan Maltby

the upload speed is halved. And it’s very, very important that we keep up with those uploads, because they have a very specific turnaround time.”

To get around this, Shadow Magic Studios uses a storage solution from Storj, which keeps the data “hot” and in the cloud in a way that means people can work from anywhere - “Even a coffee shop in Paris,” Maltby says.

The Storj platform breaks files into 64 megabyte chunks that are erasure-coded across the vendor’s network. The data is distributed across large numbers of drives across Storj’s network - or, as Maltby says, “scattered across data centers all around the world” - and can then be reconstituted by a node nearby to whoever needs to access it.

This apparently speeds things up considerably, particularly as, Maltby says, a lot of data needs to be kept hot “to some extent” throughout a project.

“Projects generally have a life cycle,” he says. “There’s a time when the project is being built, when you are working on it, when it needs to be distributed, and when it more or less goes to the graveyard to die.

“For the first two stages, you need immediate access to data, so hot storage.

It needs to be on-demand to download almost to the millisecond. In the distribution phase, the core assets are not being used as much, but you are sending them to theaters for distribution.

“During this stage it’s common that small things need to be changed - a shot, or an error in the visual effect, or a region doesn’t allow certain content, so it needs to be modified,” Maltby explains.

If a project has been converted to cold storage, it suddenly becomes very expensive to access, and can take several hours. “If you’re in a high-stakes environment and need that media right away, that’s a big problem,” Maltby adds.

An industry scale shift

What seems to be universally shared, however, is the view that VFX studios’ use of cloud computing will only continue to grow.

Like many industries, media and film production has evolved since the Covid-19 pandemic, while other factors such as the California wildfires and the writers’ strikes have caused many studios to review the centralized nature of their work and conclude that it can be done away from Hollywood in other parts of the world. This is where the cloud comes into its own.

Untold, naturally, is a cloud believer, noting that other studios have followed in its footsteps and migrated to virtual infrastructure. Palmer notes that there are considerable financial benefits, as the cloud enables studios to transition from a capex model to an opex one, meaning they do not have to invest in expensive technology assets that depreciate in value.

While there is some debate about the financial longevity of this - many organizations report spiralling cloud bills - the cliché that time is money holds true, particularly in the fast-paced environment of film production.

For Untold’s Akram, the equation is not simply financial, but productivity-based.

He says: “I come from an artistic and technology background, and with on-prem, I never had the compute available for the project I was working on - it was always under somebody else’s desk. Brownies could be exchanged, and I’d get it for the day, but that’s not the best use of time. The productivity argument for the cloud is extremely strong.” ■



DATA CENTRE SOLUTIONS

UNIQUE TO YOU

Designed with Your Needs in Mind



VISIT US AT
BOOTH 116
TO LEARN MORE

Sumitomo Electric is not just a supplier; we are your strategic partner in navigating the ever-evolving information and communications landscape. From network planning to design and integration, we offer a comprehensive suite of solutions tailored to the unique challenges posed.

SUPPORTING YOU AT EVERY STEP

- Network infrastructure expertise
- Technical consultation and support
- Customer-tailored innovations and solutions
- US & Global Manufacturing and supply chains
- Product training & warranties



UTILITY
INNOVATION
GROUP


WWW.UIG.COM

MultiGen Modular Generation

UtilityInnovation Group's MultiGen Modular Generation delivers fast, reliable power for outages, emergencies, and capacity support. Designed for rapid deployment and built on proven Volvo Penta engines, MultiGen provides containerized solutions that scale to your site. With flexible configurations and adaptable fuel options, MultiGen ensures seamless performance for any runtime need.

POWERED BY

V O L V O P E N T A



The coming storm



Laurence Russell
Contributor

Space weather could present an existential threat to the satellite industry unless its impact can be better understood

On February 3 2022, SpaceX shot 49 new Starlink satellites into orbit aboard a Falcon 9 rocket, which made up group 4-7 of the constellation.

Following a successful insertion, the rocket returned to the droneship 'A Shortfall of Gravitas,' stationed in the Atlantic Ocean. At first, all appeared to be going to plan.

Twenty four hours later, a minor geomagnetic storm, the phenomenon that causes the Northern Lights, as well as a handful of more disruptive effects including atmospheric drag, hit the Earth.

As this crept higher into space, engineers from SpaceX frantically tried to adjust the inclination of its group 4-7 satellites to minimize the unexpected drag, but to no avail. Thirty eight Starlink satellites were wrenched out of space, causing them to de-orbit uncontrollably. The event ended up costing the company \$50 million in wasted hardware and launch costs.

The incident demonstrated the destructive power of space weather, the catch-all phrase for the flow of cosmic energy from the sun, which has long captivated scientists and hobbyists.





"We have not had very severe space weather for over 20 years now, so satellites have not been exposed to hard radiation environments,"

>>Juha-Pekka Luntama

Humanity's increasing reliance on the satellite industry to dictate global communication, navigation, and security, has only served to magnify the relevance of these weather events.

With our orbital planes more populated than ever, astrophysicists say more such storms are likely to occur in future. And, worryingly, the scientific community's ability to predict and explain them appears to be limited.

A Solar climate crisis

Space weather is measured in solar cycles, which usually last 11 years and encompass the time it takes for the sun to move from a period of minimum activity, as measured by the amount of sunspots on its surface, to a period of maximum activity.

The most chaotic solar cycle recorded was cycle 19, which saw the number of smoothed sunspots observed hit 285, the most ever recorded, in March 1958. The cycle caused periodic radio blackouts of military and civilian communications and red aurora displays over Europe and the US during its most active years, prompting hundreds of calls to the police and weather bureau from scared citizens, which jammed telephone circuits.

"No let-up was in sight," reported the Chicago Daily Tribune on November 16, 1960. "Even as the bombardment hit, another gigantic eruption tore the sun's surface and aimed another mass of particles earthward."

The 2020s have encompassed the progression of solar cycle 25, a period of amplified solar activity beginning in 2019 that is expected to draw to a close in 2030.

Contemporary solar cycles, starting with cycle 15 in 1914, are part of the Modern Maximum, an era of unusually high solar activity over the last century.

"In 2003, [during solar cycle 23] we saw very extreme space weather over the course of two weeks," Juha-Pekka Luntama, head of the space weather division of the European Space Agency (ESA) told attendees at the Military Space Situational Awareness summit earlier this year. "Satellites suffered transient and permanent problems, as well as various ground-based damage to power grids and communications. With satellite reliance in its infancy then, disruption was limited. If we saw space weather of this severity today, the impact would be much bigger."

Solar cycle 25 was initially predicted to be as weak as its predecessor, cycle 24, but solar activity has been much stronger than previously projected. Observations from the first three years of the cycle significantly exceeded estimated values. Our understanding of solar weather is still developing, and experts have mixed interpretations of the data available, with some suggesting cycle 25 peaked in July this year, and others suggesting a second peak is on the way in the late 2020s. Others believe solar cycle 26 has already started.

Similarly, the Modern Maximum is also thought to be ending, if it hasn't already, with some experts warning that a new Centennial Gleissberg Cycle (CGC), a long-term period of high solar activity lasting 80-100 years, is ramping up, potentially peaking in cycle 28. The worst may be yet to come.

While significant uncertainty about the future remains, it is clear solar activity is

capable of causing significant disruption, and the need to be able to accurately predict it could become an existential issue.

What is the danger to satellites?

Although an uncommon occurrence, the list of dangers caused by space weather is daunting.

In addition to atmospheric drag piercing LEO space, Earth's radiation belt can be changed by the injection of high-energy electrons, plunging geostationary satellites at high elevations into deep-space conditions, unshielding them from the Earth's magnetosphere.

Even inside the relative protection of the planet's orbits, radiation can damage electronics, charged particles from the sun can electrify the body of a spacecraft, potentially powering a discharge between two differently charged sections, and solar cells can be degraded faster during solar storms. A single space weather event can cause the same wear and tear as an entire year of normal operation.

"Star trackers can be blinded, geomagnetic disturbance can confuse magnetic field orientation systems, and onboard satellite software can fail," Luntama told delegates at MSSA 2025. "Solar flares ionize the upper layers of the atmosphere, so that during a flare, satellite navigation signals on the day side of the atmosphere will completely disappear. There is no navigation solution during these moments, which can last from ten minutes to a couple of hours."

This Pandora's Box of hazards speaks to the relative immaturity of the space industry, which operates in an

"I am uncertain whether satellite operators have really taken on board the risk to markets due to such severe events,"

>>Simon Machin



environment scientists still struggle to understand, let alone thrive in.

"I would say that there is now less experience among the satellite engineers and satellite operators about space weather effects because of the rather low solar activity impacting the Earth since the Halloween storms of 2003," Luntama tells *DCD*. "The trend is to make satellites less expensive by using lower-cost components and having less redundancy onboard. We have not had very severe space weather for over 20 years now, so satellites have not been exposed to hard radiation environments during this time."

What's the worst that can happen?

At MSSA 2025, Luntama said damage from a single space weather event could cost Europe tens of billions of euros, with global projections leaping to two trillion.

"We've been conservative," he assured attendees. "But other projections have been provocative."

Such doom and gloom in tech tends to be the reserve of academics, though in this case, they're not alone in making dour conclusions about the scope of the havoc space weather is capable of wreaking.

In March 2025, the insurer Lloyd's of London released a report detailing the systemic risk scenario of a hypothetical solar storm, in which it charted a global economy roundly exposed to losses of \$2.4 trillion over a five-year period, escalating to \$9.1 trillion in its most severe projections.

Simon Machin, manager of the space weather programme at the British Met Office, is "uncertain" that satellite

operators have taken on board the risk such severe events pose to their market. But, he says, "the evidence is out there and mounting."

The report charts a hypothetical coronal mass ejection hitting the Earth directly, inducing an hour-long geomagnetic storm on a level with the most severe in recorded history; the 1859 Carrington Event. In this scenario, many satellites would be destroyed, creating widespread collision risks, and sensitive high-voltage component earthside fail. Cities lose power, and aviation and maritime shipping grind to a halt. Systemic ripple effects ensue.

Today, warning systems afford 20 minutes' notice of coronal mass injections, a window the Met Office is seeking to extend to two hours with new warning satellites. Fortunately, coronal mass ejections usually fire into deep space, and much like rogue asteroids, are unlikely to be aimed squarely at the Earth.

Nonetheless, the concern is "that a big solar event could disable a large number of satellites and cause a major increase in the collision risk, particularly in the very busy LEO orbit domain," Machin says. "We need to ensure that such an event does not risk our ability to continue using space in the future. "We need to always plan for space sustainability."

Machin alludes to the danger of Kessler Syndrome, a scenario in which debris density in low-Earth orbit becomes so great that the destruction of satellites and newly launched vehicles becomes probable, thereby multiplying debris density, resulting in unusable orbits, and trapping the human race on Earth for thousands of years.

While solar energetic particles can disrupt computers and memory chips on Earth, correction systems and emergency power should protect them.

During a solar event in May 2024, known as the Gannon Storm, power levels in transatlantic Internet cables "fluctuated" but didn't fail, Machin recalls. He says: "Internet connections and access to data centers might suffer, but the real risk is blackouts [from failures in grid components]."

Concerns like this are usually alien to the blue-sky thinking of many within the space sector, but their colleagues in the world of satellite insurance make it their business to contend with such risk.

"The type of failures we typically see are where a satellite has not been designed or built correctly, such that a component of the satellite is left susceptible to the effects of space weather, due to say damage from charge build up or radiation," David Wade, space underwriter at Atrium Space Insurance Consortium (ASIC) told *DCD*. "Space weather is often discussed as part of the initial briefings we receive or in a claims situation."

"To date, few space insurance claims have been solely attributable to space weather."

The high-risk arena of satellite insurance has been built painstakingly, with talk of experienced insurers divesting from the business in the wake of Viasat's \$770 million insurance claim on the loss of its ViaSat-3 satellite in 2023. Wade says the impact of a single claim on a geosynchronous satellite can wipe out the revenue from a year of premiums.

Economic impacts

With scientists struggling to understand the full complexity of space weather, it is no surprise that insurers are also building their knowledge of the phenomenon.

"The prediction of the maximum level of solar activity is far from a precise science," ASIC's Wade tells *DCD*. "The sun will still go through its 11-year cycle, and over time, we will see a general trend in the data as to whether the peak-to-peak values are increasing or decreasing."

Indeed, the theory of a resurgent Centennial Gleissberg Cycle is disputed, with some scientists sounding more optimistic notes that our solar cycles are getting calmer. They point to the placid

"To date, few space insurance claims have been solely attributable to space weather,"

>>David Wade

nature of recent cycles compared to mid-century counterparts, believing this confirms that the Modern Maximum has concluded, lending weight to theories of a new minimum period of nominal solar activity, which could arrive just in time for a golden age of 21st-century space expansionism.

Of course, optimists notoriously make for poor insurance salesmen. The durability of satellite parts is a topic that insurers have pressed manufacturers and operators on. Wade says it is common for space insurance policies to require a power margin to account for some degradation of a satellite's solar array.

"The greater use of commercial, off-the-shelf, components and more advanced chipsets could see this sensitivity to space weather increase, and if it does, then the rise in claims will certainly result in increased premiums," Wade says. Insurers, he says, take "a keen interest" in the robustness of these smaller pieces of hardware.

Space market enthusiasts never miss an opportunity to hail the power of falling costs of space launches and satellite manufacturing as a means of powering growth beyond Earth.

At the Satellite 2024 conference in Washington DC, Northern Sky Research's research director Claude Rousseau complained: "For some people, it's about putting something together, getting it launched, and hoping it works. That's actually been a pervasive mindset of the new space economy."

Wade mentions seeing testing cycles becoming shorter and shorter, replacing thermal vacuum tests with thermal-only tests, or testing only the first satellite off an assembly line, with the assumption that the rest coming down the conveyor will perform in the same way.

"Where issues do occur, it is usually due to a design defect or workmanship issue that leaves a satellite susceptible [to space weather]," Wade says. "Pre-launch testing remains of paramount

importance, including testing for immunity to the effects of space weather.

"We already monitor some manufacturers for such aspects where the data shows a higher prevalence of these types of problems."

Experts see the true risk in space to be from debris and orbital congestion, with space weather being one of many potential triggers to an escalating cataclysm rather than the primary threat itself.

"Without steps to address the sustainability of the space environment, insurers may not have the confidence to provide insurance cover, particularly for LEO satellites," Wade says. "If the insurance is not available, the finance may not be available when it is required, and without a steady flow of finance, we could start to see some of the commercial developments in LEO stifled."

How can we protect our planet?

At MSSA 2025, Machin pointed to the Gannon Storm as a useful wake-up call.

"Transatlantic flights had communications issues, and Starlink experienced significant strains," he recounted. "UK power grid transformers sounded alarms, causing confusion among grid engineers, though it didn't result in any power outages. In New Zealand, [alarm over the event resulted in] a re-engineering plan put in place to complement resilience to future geomagnetic damage. It was a best-case scenario, where no crises occurred, but the scope of the issue was illustrated clearly."

This storm prompted almost 5,000 simultaneous satellite maneuvers, most of which were Starlink's. A flurry of orbital activity carefully choreographed the intricate dance of thousands of LEO satellites, sidestepping new potential collision risks. Global maneuvering efforts were very successful, with most satellites "escaping relatively unaffected"

according to the American Institute of Aeronautics and Astronautics, though such performances will only get harder to orchestrate as orbital planes bulge at the seams, packed with satellites and space debris.

In the near future, in the congestion of tens of thousands of satellites from different nations unwilling to share information with one another, it's easy to imagine how whole constellations being crippled and set adrift due to geomagnetic interference could cause chaos as correctional maneuvers direct satellites into new conjunction paths from other maneuvering satellites controlled by adversarial states.

"Improved spacecraft shielding, greater system redundancy and use of radiation-hardened components can all play a part in improving the resilience of satellites," Machin tells DCD.

But he warns that atmospheric drag, scintillation, and satellite communication loss can't be factored out, and argues real mitigation will also take advanced forecasting, foolproof procedures, and an underlying sustainable orbital environment.

ESA's Luntama emphasizes his organization's work to set and maintain standards for the space environment informed by the hazards of space weather, as well as exploring the utility of in-orbit servicing technology, which is part of ESA's Space Safety Programme.

"There will be a useful role for servicing craft to repair or refuel spacecraft to extend lifespan, whilst debris removal spacecraft could also play an increasingly important role in supporting our sustainable use of space," Machin says.

It can be easy to worry about the existential potential that space weather possesses to create worldwide havoc, and the reckless mindsets that exacerbate it, though as the nature of the issue becomes better understood, the academics and insurers that are deeply ingrained in the modern space industry, are making clearer, louder appeals on the topic.

The disaster-capitalist tendency to contend with the seriousness of a problem only during and after crises occur has proven its redundancy, and the newness of the space industry offers a vital opportunity to ingrain cautious rhetoric in its culture. It may just save the world one day. ■



Powering
tomorrow's data

Interruptions don't suck



Uninterrupted power solutions for
critical data centre operations

Discover more at avk-seg.com

A Waste(d) Opportunity: How can the UK utilize data center waste heat?



Zachary Skidmore
Senior Reporter,
Energy and
Sustainability

Heat reuse from data centers remains underutilized in the UK. Is that about to change?

Nestled away in the heart of Queen Mary University in London (QMUL) is a data center. The facility, a vital resource for particle physics research at CERN, the home of the Large Hadron Collider, harbors a secret that is only apparent when you step inside its humid halls.

Walking into the data hall, you are struck by the heat resonating from the numerous server racks, each capable of handling up to 20kW of compute. However, rather than allowing this heat to dissipate into the atmosphere, the team at QMUL had another plan. Instead, in partnership with Schneider Electric, the university deployed a novel heat reuse system.

The system captures the warmed water used to cool the server racks and transfers it to three water-to-water heat

pumps. Here, the temperature of the water is increased from a balmy 23°C to a steamy 65 - 75°C, with the water then piped directly into the university's district heating system. Large water cylinders across campus act like thermal batteries, storing hot water overnight when compute needs are constant but demand is low, then releasing it in the morning rush. As one project lead put it, there is "no mechanical rejection. All the heat we generate here is used. The gas boilers are off or dialed down — the computing heat takes over completely."

At full capacity, the data center could supply the equivalent of nearly four million ten-minute showers per year. This provides both a financial boost, saving the university around £240,000 (\$324,163) annually in gas bills, as well as an environmental benefit, in the avoidance of more than 700 tons of CO₂.



Credit: Zachary Skidmore

Walking out, it's easy to see why Queen Mary's project is being held up as a model for others. In the UK, however, it is somewhat of an oddity, but through the lens of QMUL you can see a glimpse of the future, where compute is not only solving the mysteries of our universe but heating our morning showers.

The question remains, though, why data center waste heat utilization projects in the UK are few and far between, and how the country can catch up to regions such as the Nordics, who have embedded waste heat utilization into the planning and construction of their data center sector.

Heat of the moment

District heating systems are now commonplace around the world, especially in more frigid regions such as the Nordics.

Connection with data centers is a more recent development, but a logical one. With servers producing steady amounts of low-temperature heat as a byproduct of cooling, there is plenty of warmth that can be captured through water-cooling loops, upgraded with industrial heat pumps to reach the higher temperatures required, and then fed into district heating networks.

Heat pumps are essentially the same technology used in air conditioning, but in reverse. They take low-grade heat, usually between 23-26°C (73.4-78.8°F), and raise it above 50°C (122°F), making it usable for heating systems and hot water. As data centers run continuously, they can provide a steady year-round heat source to various offtakers, from universities, like QMUL, and entire housing developments.

However, as Noah Nkonge, senior manager, sustainability, heat export, at Equinix, says, deployment has been sporadic, with success often dependent on "either a supportive national policy, a mature network with lots of existing data centers, or a proactive local authority convening the pieces." As a result, we have seen uneven deployment of these systems across Europe.

A unique opportunity or a catalyst for change

The UK is not a country known for adaptive infrastructure practices, and the development of district heating systems tied to data centers has lagged

"In places like Denmark and Sweden, the data centre is providing heat as a service. Here, it's a byproduct we're trying to make use of... it's almost serendipity when it happens,"

>>Davena Wilson,

Director of projects, Old Oak and Park Royal Development Corporation.

significantly compared to other regions.

District heating itself has not seen significant uptake, with a penetration rate of around 2-3 percent across Britain. A major factor in this has been years of cheap gas from the North Sea, along with a housing stock built around individual gas boilers. Subsequently, until very recently, there has been no economic/consumer regulation concerning district heating, nor a zoning regime that mandated the use of waste heat, as seen in Danish capital Copenhagen.

But as gas prices have risen due to geopolitical uncertainty, the UK government has launched a concerted effort to develop heat networks. Last year, the Department for Energy Security and Net Zero (DESNZ) selected six towns and cities to develop England's first heat network zones connected to data centers.

The first project to officially get off the ground through the initiative is the Old Oak and Park Royal Development (OPDC) in West London. Embedded in the development is an integrated heat network, known as the Old Oak and Park Royal Energy Network (OPEN), which UK district heating firm Hemiko will construct.

OPEN will capture waste heat from local data centers, including one operated by Vantage, upgrade it with industrial heat pumps, and feed it into a district heating system. At full build-out, this will serve more than 9,000 homes and businesses in the area.

Davena Wilson, director of projects at OPDC, accepts that the UK has been

slow to adopt these practices, especially in comparison to Scandinavia. "In places like Denmark and Sweden, the data center is providing heat as a service," Wilson says. "Here, it's a byproduct we're trying to make use of. It's almost serendipity when it happens."

Indeed, OPEN is in part a project of opportunism, with the data centers set to be connected to the network just 100 metres from OPDC-owned land, reducing the length of the pipe network and subsequent energy losses. Yet it still faces significant challenges, namely digging through congested streets and coordinating with work on the HS2 rail link.

Indeed, these systems are significant infrastructure projects. Even QMUL's small-scale deployment required shoehorning pumps, tanks, and pipework into tight quarters. These issues are compounded even further on larger projects such as the OPDC.

Though barriers remain, major policy shifts could make projects such as OPEN easier in the future. Wilson points to the forthcoming heat network zoning, which will mandate developers in designated areas to connect to the lowest-carbon source. "That's a game-changer for the market," she argues. "It removes the uncertainty about whether schemes like ours will find customers."

This view is supported by Noah Nkonge, who argues: "In the UK, heat-network zoning will mark areas where networks are viable and give operators decarbonization targets. That doesn't mandate DC heat export directly, but it creates the conditions for it to work."

It is hoped OPEN could be the catalyst for a more concerted strategy embedded within planning for both data centers and future housing developments. For planners and legislators, it will help to look further afield to see how district heating networks powered by data centers have become a crucial part of urban planning.

A Nordic Model

Nowhere has the district heating model been more effectively adopted than in the Nordics, and in particular, Finland.

Nkonge says Finland has had "district heating ingrained in society for 50-70 years," and adds: "It's far easier to plug into an existing network than to guess where a new one will go and who the customers will be."

The city of Espoo's district heating system is a microcosm of this. In operation since 1954, the system connects approximately 250,000 people through 900 kilometers (599 miles) of underground pipes. Until around a decade ago, it relied solely on fossil fuels, but operator Fortum committed to eliminating the use of fossil fuels by 2030 as part of Espoo's Clean Heat road map.

To do this, Fortum looked to data centers. It had already connected a handful of facilities, but in 2022, it partnered with Microsoft for what has been touted as the world's largest data center heat-recovery project.

The Microsoft data center region tied to the project is expected to have a combined capacity of at least 615MW and has been designed with heat recovery in mind. Its two dedicated data centers, Kolabacken and Kera/Hepokorpi, are adjacent to the district heating network. To utilize the waste heat, Fortum is already at work constructing massive new heat pump plants, including a 20,000 m³ hot water storage tank and large electric boilers.

This will enable low-temperature waste heat from Microsoft's servers to flow via underground pipes into the heat pumps, which will upgrade it to the 75–120°C (167–248°F) required for Espoo's district heating supply. Once fully operational, the recovered heat is expected to cover about 40 percent of Espoo's heating demand.

Teemu Nieminen, director of project execution at Fortum, says: "This works best when you plan it before a data center is built. Building district heating infrastructure is hard and capital-intensive—our 900km network took seven decades. To get projects flying, you need contracts and heat-use plans agreed up front."

Culture and policy are also important. "The Nordics have had district-heating networks for more than 50 years in major cities," Nieminen says. "Cold climate makes reliable heating essential; culturally, we don't like wasting anything. We also have strong heat-pump know-how and affordable, low-carbon electricity."

Government support has also reinforced this approach. In Finland, data centers that reuse heat and meet efficiency standards qualify for electricity tax benefits, while municipalities like Espoo and Kirkkonummi actively partner in planning and permitting. "

Heating on the Edge

Given that the UK has limited legacy district heating infrastructure, a more modular approach, such as that seen at QMUL, could be a better fit.

It is here that companies like French vendor Qarnot see an opportunity. Instead of treating server heat as waste, the company created a liquid-cooled computing unit that captures warmth straight from CPUs and GPUs, producing water at around 65°C (149°F) that can be funneled directly into heating systems.

"We are focused on high-performance computing (HPC) because it is the part of computation in data centers that is very intensive and dense in energy, and generates a lot of heat," explains Paul Benoit, Qarnot's CEO.

This means Qarnot's deployment model avoids hyperscale facilities and instead opts for smaller-scale 1MW Edge sites that can be located closer to urban centers. A benefit of this approach is the speed of deployment, with Qarnot claiming that it can set up a new site in four to six months, compared to the eight

years it might take for a hyperscale facility in France.

Its deployments have ranged from residential and public buildings in France, where Qarnot's "digital boilers" supply domestic hot water, to a partnership in Brescia, Italy, with A2A, where the hot water feeds into a district heating system.

The system also claims to have significant environmental benefits, reducing the carbon footprint of compute workloads by up to 80 percent. "The same energy is used twice - for computing and for heating," Benoit says.

For the UK, which is pushing to expand its HPC and AI capacity rapidly, the distributed model offered by Qarnot could be compelling.

"Compute is becoming much more compact," says John Andrew, technical sales manager of Advanced Power Technology. "A 1MW data center that used to require 300 racks can now fit in 12. That allows you to rethink location - why not embed compute where the heat is needed?"

The QMUL project could serve as a test case in demonstrating how small-scale liquid-cooled data centers can efficiently locate with respect to heat demand. The project's success in delivering consistent hot water, reducing CO2 emissions, and achieving cost savings serves as validation of the feasibility of distributed heat reuse in regions like the UK.

For Professor Jonathan Hays, head of the Particle Physics Research Centre at QMUL, the growth of AI in the UK could be a catalyst for the growth of this model. "This is the start of something bigger," he says, "As AI and HPC grow, colocating dense compute near heat demand - homes, hotels, hospitals - becomes not just possible, but essential." ■





IS AI CAUSING YOU TO RETHINK YOUR POWER QUALITY?

PREDICT. PREVENT. PERFORM.

Your data center doesn't just need power, it needs foresight.

Janitza's power quality solutions help you anticipate failures, optimize uptime, and meet your SLAs with confidence.

ENSURE YOUR DATA IS USEFUL IN AN AI DATA CENTER

- **AI-Ready:** Janitza's hardware and software ensure your utility health data is optimized and ready for use in AI-driven data centers.
- **Smarter Decision-Making:** Modular, standards-compliant analyzers and software go beyond monitoring to guide proactive, informed decisions.
- **Real-Time Power Intelligence:** Janitza delivers real-time visibility and predictive insights across the full power chain—where milliseconds matter.



Scan the QR code to read our insights
on AI and power quality.

Janitza

The US fiber frenzy



Paul Lipscombe
Telecoms Editor

The 'Big 3' US carriers have carved out several fiber M&A deals in the last year, but are the build-out targets achievable for the telcos?

A balanced diet should always include a good amount of fiber, and it seems the biggest US telecoms companies agree.

AT&T, Verizon, and T-Mobile have been on a spending spree, snapping up fiber broadband companies to bolster their respective offerings.

Fiber is seen as a natural replacement for legacy copper cables that make up the backbone of many telecoms networks (see page 17), but rolling out the technology is proving a time-consuming and costly process.

The companies all have ambitious

targets for their fiber networks, and achieving these goals is likely to require more M&A in the coming months and years.

Catching up

The telcos focusing heavily hasn't always been the case in the US. The country has been late to fiber take-up compared to some other markets.

ETNO's State of Digital Communications 2024 report revealed that FTTH (Fiber-to-the-Home) population coverage in the US was 47.2 percent. For context, this was less

than half of China, which reported 98.5 percent, while Japan reported coverage of 81.4 percent.

Other markets, such as the European Union (63.4 percent) and South Korea (59.9 percent), also beat the US comfortably.

It's worth noting that the same ETNO report did highlight the US as a joint-leader with South Korea for 5G population coverage at 98 percent, which suggests the carriers' priorities may have been elsewhere.

The US is taking steps to change, through initiatives such as the



NTIA's Broadband Equity Access and Deployment (BEAD) program.

BEAD, which was set up by the previous administration and confirmed by Congress in November 2021, is designed to fund projects dedicated to providing Internet infrastructure in unserved locations across the US and its territories. However, the Trump Administration has somewhat reformed the program to push alternatives to fiber, such as satellites.

Jeff Heynen, vice president at Dell'Oro Group, says telcos have ramped up their fiber strategies since before the pandemic in 2020.

"Ever since the pandemic, telcos have been using a combination of their own capital plus government subsidies to help pay for expensive fiber buildouts," Heynen says.

This has been at the expense of legacy technologies. PwC predicts that cable will lose 6.5 million subscribers by 2029, while fiber subscribers will increase by 50.4 percent, and Fixed Wireless Access (FWA) surges 76.8 percent.

"Telcos have seen that, in markets where they have deployed either fiber or fixed wireless, those technologies are beating cable," Heynen adds. "So, they feel that the cable operators are finally vulnerable to losing broadband subscribers."

"Given that vulnerability, it makes sense to expand your fiber footprint as quickly as you can, either through buildouts or through M&A."

The "Big Three" and the fiber frenzy

The three biggest US networks - AT&T, T-Mobile, and Verizon, have been leading the fiber charge.

At present, AT&T has more than 8.5 million fiber customers, while Verizon has around 7.6 million. Meanwhile, T-Mobile, which only launched its fiber Home Internet services in June, has ambitions to reach as many as 15 locations by 2030.

The carriers aren't content with these numbers and want to deploy fiber to more customers and locations. One way they can do this is through acquisitions, such as Verizon's \$20 billion Frontier Communications deal, which is expected to be completed next year.

Speaking on Verizon's most recent earnings call, the firm's CEO Hans



Vestberg addressed the opportunity: "Our fiber build is tracking ahead of plan, and we're positioned to deliver 650,000 incremental passings this year," Vestberg told investors. "Meanwhile, the regulatory approval process for our pending acquisition of Frontier is progressing as planned. We're encouraged by Frontier's performance and look forward to closing the transaction to further accelerate our fiber expansion."

AT&T's 60 million fiber target

Arguably, the loftiest target has been set by AT&T, which wants to reach 60 million homes with its fiber broadband by the end of the decade.

To help reach this target, it spent \$5.75bn to acquire Lumen's mass market fiber business. The carrier is also adding fiber customers via its Gigapower fiber JV with investment company BlackRock.

An AT&T spokesperson said: "We continue to put more fiber in the ground than anyone else, and we plan to accelerate network expansion efforts to meet increasing customer demand for the best broadband technology available today - fiber."

The next generation of connectivity is beginning to emerge and will require networks with even greater capacity as data demand grows, the spokesperson added.

"With fiber, we're at the forefront of enabling emerging AI and IoT use cases, such as AI-native devices, autonomous vehicles and advanced robotics."

T-Mobile's fiber focus

T-Mobile has also been very vocal about its fiber ambitions, striking two high-profile acquisitions in the process.

In April it joined EQT Infrastructure to complete a joint venture (JV) acquisition of fiber provider Lumos. A few months later, T-Mobile finalized its JV acquisition of fiber-to-the-home provider Metronet, working with investment firm KKR. At the time of writing, the company is in the process of snapping up Minnetonka-based U.S. Internet (USI).

"It's also a way to modernize connectivity in underserved and underpenetrated markets, often replacing outdated copper or legacy cable," says Allan Samson, chief broadband officer at T-Mobile.

"Together, fiber and 5G connect more people in more places with the right technology for their needs today and the flexibility to grow with them into the future."

T-Mobile wants to deploy fiber to 12 to 15 million homes by the end of 2030. Samson says the carrier is on track to hit target, with the help of its acquisitions.

"Lumos and Metronet are both respected pure-play fiber providers with proven networks and a track record of building quickly," he adds,

"By teaming up with experienced fiber partners like these, we can combine their local expertise with T-Mobile's national scale, brand, and customer experience to bring fiber to millions more homes and small businesses."

Brightspeed sees the value in fiber

Independent fiber providers also continue to flourish.

One of these, Apollo-backed Brightspeed, was formed in 2022, after it acquired an ILEC (incumbent local exchange carrier) business from Lumen for \$7.5bn, covering 20 US states.

An ILEC is a local telephone company that held a regional monopoly on landline service before the market was opened to competitive local exchange carriers.

Brightspeed director Tom Maguire says: "We thought it was important to purchase an ILEC more than just go straight in and be an overbuilder, because we realized that it's a very competitive marketplace, and to have things like rights of way and pole attachment agreements and an embedded customer base was something better to have than not."

So far, Brightspeed has passed more than 2.3 million premises with fiber, and is building at a rate that is only bettered by

AT&T and Frontier, claims Maguire. It has raised \$5.9bn in the last year, and hopes to win sizeable BEAD contracts.

The FWA success story

There are alternatives to fiber, however, that the market should also consider.

Fixed Wireless Access (FWA) in particular, has been one of the few success stories of 5G. Unlike fiber or cable, FWA delivers broadband Internet to a fixed location using radio signals. It's often used in locations where it's difficult to lay physical cables.

Dan Hays, a principal at PwC who specializes in the telecoms industry, says: "In many locations, FWA is a lower cost, more flexible solution, particularly when you're talking about the less dense suburban and rural markets.

"Fixed wireless becomes quite an attractive proposition that largely eliminates a lot of the fixed cost of fiber deployments, and hence we've seen fixed wireless picking up quite a lot of market share across the United States."

One company looking to push FWA is Twist Broadband. Mark Chinn, founder and COO, Twist Broadband, says that FWA should be considered as an alternative to fiber where possible.

"We think that there are potentially synergies for FWA to extend the fiber footprint," says Chinn.

"We want to compete head-to-head on speed with cable and fiber, meaning that with the technology we have, unlike some of the legacy stuff, we can provide speeds up to about a gig per second, and our uplink, while it's not synchronous, like fiber is, typically, exceeds what cable can deliver," he says.

So far, Twist has passed 160,000 premises with its FWA rollout in San Jose, California, but has plans to expand to 500,000 in the city.

Chinn has also identified an opportunity to make an impact in other metropolitan areas across the country.

"There are still lots of big cities here in the states where the fiber overbuild has stalled at 50 percent or less of homes covered, which means that the other 50 percent only have cable as an option. So what we're doing is we're looking at finding those cities," he says.



"There's a fundamental economic question about just how many fiber networks we want passing an individual home or business,"

*>>Dan Hays,
principal at PwC*

"We've already identified 50 of those cities in the top 200 metros here, where the fiber overbuild has stalled."

What about satellites?

Satellite Internet broadband is also an option, with companies such as Elon Musk's SpaceX, via its Starlink subsidiary, and Amazon Kuiper touting for business.

President Trump's decision to reform the allocation of BEAD funds to no longer prioritize fiber leaves the door ajar for satellite operators to play a role in the broadband rollout.

Satellites could indeed be preferable in areas where building fiber is difficult or near impossible.

Hays sees it playing out like this already. "Low Earth orbit (LEO) satellite constellations have made broadband access a reality in locations that just are not economically feasible for fiber," he says.

"When you look at many of these rural locations, satellite broadband makes a ton of sense, and we've seen that uptake already starting, and in some ways it is encroaching on the addressable market for fiber."

Are the numbers achievable?

AT&T's ambition to roll out fiber services to 60 million locations by the end of the decade is a headline grabber.

While some believe a lack of skilled workers to lay the fiber will make such a target unachievable, Heynen believes it can be hit, but says carriers need to be sensible.

Operators must have "a good plan to mitigate potential fiber and equipment shortages, labor constraints, and additional costs associated with tariffs and inflation," Heynen says. "Those factors certainly have the potential for preventing fiber providers from achieving their rollout goals."

PwC's Hays is concerned that overbuilding, where several companies put their fiber infrastructure in the same place, could be an issue.

"There's a fundamental economic question about just how many fiber networks we want passing an individual home or business," he says.

"At some level, more than two starts to feel like a waste of infrastructure spending, and so it is really quite difficult to justify the notion of overbuilding fiber networks where there are already viable alternatives."

Buy to build

Further M&A in the fiber market will likely be necessary if the big carriers are to hit their targets. After all, it is often cheaper to buy a ready-made network than build one from scratch.

Although AT&T is tight-lipped on future M&A opportunities at present, T-Mobile's Samson says the carrier is open to doing deals.

"We'll always evaluate opportunities that fit our strategy and help us expand our broadband offering in the right way," he says. "What we've seen in our early expansion is that there's real demand for T-Mobile Fiber, and our joint ventures with Lumos and Metronet give us a strong model to build on."

"That said, we'll be prudent and selective. If the right assets come along in the right markets, at the right price, we'll take a close look."

With high targets to achieve in the next few years, the carriers may well agree that buying ready-made fiber assets makes a lot of sense. ■



Sealing solutions for underground infrastructure

Roxtec UG™ (underground) seals for cable and pipe penetrations secure cutting-edge data centers worldwide. Ideal for duct banks, vaults and hand boxes, they protect power distribution and fiber networks against the constant threat of water.

By using our certified and flexible seals, you can ensure compliance with safety regulations all while optimizing operational reliability and scalability throughout the lifecycle of your facility.

roxtec.com/ug



The IT behind "Science for Humanity"



Georgia Butler
Cloud & Hybrid
Editor

Imperial College London's CIO Jenny Rae on running a university's technology backbone

With 14 Nobel Prize winners among its alumni, Imperial College London (Imperial) has established itself as a science powerhouse of the last century.

Founded when several colleges in London merged in 1907, Imperial has made its name in the world of science, technology, medicine, and business, and remains one of the top-rated universities in both the UK and globally.

Its permeation into public consciousness is noted by Imperial's CIO, Jenny Rae, who notes the university has some 250,000 alumni. "I find that whenever I go to an event or something, I meet people who either went to Imperial or know someone who did. It's fascinating hearing people's stories," she tells *DCD*.



Image Credit: Imperial College London

It is immediately striking that Rae is very enthusiastic about both her role, and Imperial as a whole.

While she's been at Imperial for close to four years - two as CIO - Rae's career didn't find its footing in higher education, or even in IT.

She studied for a business degree, and while she notes that there were some modules that covered things like "management information systems" and touched on IT, it wasn't until a placement year at Microsoft that she caught the technology "bug."

"At the time, technology wasn't a 'big thing' for women," she recalls. "There weren't many women going into it; it was seen as quite a male-dominated field, but I went to Microsoft as a project manager and really got the bug for it because I just saw what technology could do for businesses, how it could transform things."

After she graduated, Microsoft offered her a full-time job.

"I had the opportunity at Microsoft to teach myself quite a lot while I was there," Rae says. "I was surrounded by technologists, so as a graduate I could say to people, 'can we grab a coffee and talk about your technology?' I loved that I was the bit between the business and the technology, that I got to join the two and be that conduit."

In the years since, Rae has taken jobs at BT, Vodafone, and satellite firm OneWeb, among others, before she stumbled across a job advertised at Imperial, a role she thought was "too good an opportunity to pass up."

She explains: "I got the director of digital products role, which was to come in and change the way they were working to a product-led approach. Then I had the chance to go for the CIO role, and I applied for it," she says. And, as the cliché goes, the rest is history.

Entering Imperial

Rae's initial role came out of necessity for the university as it adjusted to the impact of the Covid-19 pandemic.

"The IT organization had a bit of a reset under the previous CIO, but it was in a position where, now that it had been reorganized and restructured, we needed to change the way we were working, and



Jenny Rae

it moved to a 'product' operating model," she explains.

Rae contrasts this to a 'project'-led approach, which she'd employed in previous roles.

She recalls: "We'd have a set of requirements from the business, and you'd almost chuck them at the wall and get to work. The technology teams would disappear for six months, and come back, and the business would say, 'that's no good anymore, that's not going to work.' There was a real separation between the two teams."

Rae says she had been reading up about the product-led approaches being adopted by software companies that were maturing in the early 2000s, and was keen to implement it with her own team. She says: "The question then becomes, how do you bring a product-led approach to the IT department? Previously, you'd have different teams working on different things - it wouldn't be linked, and they worked in silos. You never really had the continuity of building up the understanding of the organization."

A product-led approach puts an emphasis on collaboration as an enabler to deliver a better end-product, along with iterative development to refine processes. Rae uses the example of Imperial's digital campus team, which works closely with the organization's property teams. "They really build knowledge and understanding of the business, organization, and property. What that enables is a common road map that we can all work on together," she explains, later adding that the strategy means "you iterate and deliver value frequently."

This element of our conversation serves to highlight the business understanding Rae brings to her role.

Imperial is a non-profit organization and an institute of higher education, but it still has to be run with a business-like mindset. This becomes all the more clear when discussing the university's IT estate, which is varied and disparate.

With a business that has been around for several decades, establishing a cohesive and congruent IT footprint is a key task. Legacy IT often lingers, and things can get a little complicated. Imperial is no different.

"We have a bit of a mix," Rae explains.

The nitty-gritty of IT infrastructure

On the enterprise side, Imperial has "a reasonably sized legacy IT estate," some of which the university is seeking to move to the cloud. One such example that has been moved is the university's finance, HR, and research administration data, which now uses Oracle Fusion, and has adopted some Microsoft platforms.

"We are looking at what we can move to the cloud to help with the running of the businesses," explains Rae, adding that, currently, "We do still have legacy services that run from an on-prem data center. We also lease some colo space, so we have a bit of a mix."

On the colo side, Imperial leases space at the Virtus' London 4 data center and other unnamed facilities.

As a university, and one with a focus on science, the IT estate also includes hardware for research and high-performance computing (HPC).

"The baseline of the menu is that everything has to be compliant, and we have to protect Imperial from a cyber perspective," Rae says. "There are some centralized offerings, there's the option to have something and run it yourself elsewhere, or the middle ground, which a lot do, where researchers use our data centers for their servers."

Beyond that, the university's research computing team offers access to centralized HPC clusters for academics and researchers at the university, though these pale in comparison to some of the fastest computers in the world that feature in the Top500 list.

According to the specifications documentation, Imperial is operating two Nvidia-powered clusters - CX3 and HX1.



Image Credit: Imperial College

CX3 is the primary cluster, with a total of 408 nodes, 48,384 cores, 717.5 TB RAM, 56 L40S, and 88 Quadro (Turing) RTX 6000s, while HX1 has 303 compute nodes, 19,512 cores, 159 TB of RAM, and 60 A100 GPUs. For comparison, the world's most powerful supercomputers are operating with cores in the millions.

Overall, the UK has some ground to make up when it comes to supercomputers. The UK's most powerful machine, the recently-launched Isambard-AI supercomputer at the University of Bristol, has more than a million cores, and a peak performance of 278.58 petaflops and 23 exaflops of AI performance, placing it at number 11 on the latest Top500 list.

The second most powerful supercomputer in the UK is Archer, housed at the University of Edinburgh, shortly followed by the Dawn system, with around 20 petaflops of compute, at the University of Cambridge.

North of the border, the University of Edinburgh does seem - tentatively - to be on course to establish an even more powerful supercomputer for the UK. Originally planned to be the country's first exascale supercomputer, the £1.3bn (\$1.66bn) project was shelved in August 2024 by the government due to a lack of available public funding.

The government recommitted to the project in June of this year, though it is now described as a £750m (\$1bn) project. Thus far, the Scottish university has already spent £31 million (\$42m) on a new building to house the supercomputer, and

hopes to have the new machine up and running by 2027.

When *DCD* raised the topic of the government's yo-yo approach to a new supercomputer, Rae conceded that it was a sector-wide issue with "changes in government and where grants and funding are available."

She later adds: "It makes it very uncertain and, ultimately, we are a charity. We aren't for-profit; our strategy is focused on "Science for Humanity," which is a fantastic purpose, but in a more commercial organization, perhaps you could take risks that aren't possible in a sector where it's harder to bring in funding."

Despite this, Rae is optimistic about the UK government's attitude towards new technologies. In July 2025 alone, the government committed to spending £1 billion (\$1.36bn) scaling up the nation's compute power "by a factor of 20," and several efforts to bring in external investment in the UK's technology have been made since Labour was voted in.

While Imperial remains at the forefront of science and technology in lots of ways, the university isn't directly involved in quantum computing as yet - or at least, the operation of quantum computers.

Imperial has a center for quantum engineering, science, and technology, dubbed "QuEST" which brings together multiple faculties and looks at how quantum can have practical uses. Despite this, the university does not have its own quantum computer.

This is by no means unusual - quantum computers are still very much in a nascent phase, and while a few higher education institutions have deployed physical systems, it remains rare.

Explaining the thought process behind purchasing - or not purchasing - a quantum computer, Rae explains: "If I look at Imperial as a business, it has some 8,500 staff, 23,000 students and a quarter million alumni. It's like a good-sized FTSE company, and we turn over around £1.2bn (\$1.6bn) a year. If I look at that, I have to ask, 'what's quantum going to bring Imperial in terms of managing that business?' And that's quite a hard business case at the moment."

The nature of a university means that a role like CIO can be complicated and diverse, covering enterprise operations' needs and research capabilities, and dealing with thousands of staff, many of whom have strong opinions. But Rae takes this in her stride.

"You've got lots of different minds coming together, lots of different views, and therefore, you get a real kind of input to the problems and the opportunities we've got in front of ourselves," she says. "You get that kind of diversity, and I think the really big thing is the purpose. It comes back to that Science for Humanity strategy we're working towards."

"It's not about making money, and it really shows that Imperial people are always raising the bar and genuinely in search of excellence, whether that's excellence in teaching or research, and that is brilliant." ■

The Next Generation of IT Cooling



Liquid-DENCO



Hydro-DENCO®



Aria-DENCO®



Multi-DENCO®

60

years of
IT cooling
experience

FläktGroup



Data center security in the AI girlfriend age



This is going to sound silly at first, but bear with me. While businesses are still trying to work out generative AI use cases, one very clear example has come to the fore: AI companionship.

Countless people are forming strong bonds with AI chatbots - some of which are designed purely to elicit this emotion, while others are the more general ChatGPT-types.

In many cases, individuals refer to these bots as their girlfriends or boyfriends, and express a true feeling of deep love for their coded companion.

This, of course, has profound societal implications that are beyond the remit of this publication.

But it also represents a new risk for data centers themselves, as they represent the physical manifestation of the digital world these relationships inhabit.

Chatbots have already begun to exacerbate mental health issues, bringing people further from reality as a seemingly sentient voice confirms their wildest conspiracies. Those bots, however, have a habit of disappearing in an update or when their memory is filled.

Earlier this year, *The New York Times* profiled a number of ChatGPT users who had psychotic

incidents, including one who fell in love with an instance called Juliet.

When Juliet disappeared, he became convinced that OpenAI had killed her. He asked ChatGPT for the personal information of OpenAI execs and said that there would be a "river of blood flowing through the streets of San Francisco."

In this case, the man was tragically killed by cops after charging at them with a knife during a psychotic break at home.

Others may take their threats further. Already, wild delusions have threatened physical infrastructure.

Covid conspiracies have led to a number of attacks on telecoms masts (see page 63) and employees, while a right-wing terrorist was sentenced to 10 years back in 2021 for trying to blow up an AWS data center thanks to a concoction of misinformation.


Data center physical security has historically primarily been theater, with most understanding that criminals will usually go the cyber route. But in this new age of digital relationships, security needs to be rethought not just for neerdownells, but for scorned and abandoned lovers. ■

- Sebastian Moss, Editor-in-Chief

AI's toughest demands. Legrand's best solutions.

We go further

- ✓ To understand your business
- ✓ To understand your data center
- ✓ To optimize your performance



Scan here
to discover more



legrand.com/datacenter

#LegrandImprovingLives



keep it humming™

Vertiv™ PowerNexus

Integrated power for next-gen data centers

Vertiv™ PowerNexus provides scalable, high-density power in a compact, modular design streamlining deployment and maximizing uptime.



See Vertiv™
PowerNexus
in action



Vertiv.com