CONTINUITY PLANNING

Five ways to build resilient data centres

Remote storage facilities are becoming increasingly vulnerable to extreme weather events. Can improving their ability to withstand such threats be achieved in an environmentally sustainable way?

Natasha Khullar Relph

On 19 July, as the UK faced record high temperatures, Google Cloud's data centres in London were experiencing cooling failures, resulting in connectivity problems and outages. Oracle's data centre was also forced into a protective shutdown, owing to what the company called the "unseasonably high temperatures".

As global temperatures continue to rise, the changing climate threatens the uninterrupted services of data centres. In a recent survey of operators by the Uptime Institute, 45% of respondents reported that they had experienced an extreme weather event that had threatened the continuous operation of their facilities. Moreover, 9% confirmed that they had suffered an outage or significant disruption as a result, which made extreme weather one of the biggest causes of service failures.

The number of data centre outages around the globe is increasing year on year, although this is because more centres are being built than ever before. According to the International Data Corporation, about 500,000 centres were handling the world's data traffic in 2012. The total in existence today is close to 8 million.

"The industry is getting much bigger and certain companies in it are becoming more powerful," observes Andy Lawrence, executive director of research at the Uptime Institute. "When they fail, more fails."

He notes that we're all becoming far more dependent on data centres. This means that, when one does fail, it has a wider-reaching impact. A quarter of respondents to the Uptime Institute's survey said that their most recent outage had cost more than $1m (£900,000) in direct and indirect costs, with a further 45% reporting that theirs had cost them between $100,000 and $1m.

Data centres are notoriously bad for the environment. They have the same-sized carbon footprint as that of the aviation industry and are set to account for 3.2% of the planet's total greenhouse gas emissions by 2025, while consuming a fifth of the world's electricity.

Consequently, efforts are focused on how data centres can meet the demands of digitisation and create infrastructure resilience, while having as little impact as possible on the environment. These are five of the most popular solutions.

1

Waste heat utilisation

Across Europe, tech companies are experimenting with waste heat recovery from their data centres. Meta has been reusing heat from its centres to warm 6,900 homes in Denmark, for instance. Microsoft, meanwhile, has powered a data centre in Finland with carbon-free energy and recycled the waste heat to nearby homes and businesses. Energy-efficiency agency Codema has partnered with an Amazon data centre in Ireland to capture waste heat for use in homes and council buildings. And in Sweden a project called Stockholm Data Parks has been running in partnership with the city's government, the local heating and cooling agency, and several data centres. The goal is to heat 10% of the capital by 2035.

"In Germany, data centres have evolved from being enemies of the state to becoming one of the heat sources," reports Stefan Mink, head of TechOps hosting at Ionos, who has been responsible for the planning, construction and management of 20 data centres in Europe and the US. "It's become a circular economy, whereby the data centres are using the energy but then also providing energy in terms of heat use."

A 2017 white paper from the Alliance for the Strengthening of Digital Infrastructures in Germany had noted that the 13 billion kilowatt-hours of electricity that was converted into heat in the nation's data centres over the year would, if reused, have met the annual energy needs of Berlin.

In 2019, investment analysis of waste heat from data centres showed that the process of reuse was a financially viable option and could provide an attractive return on investment for companies. Moreover, by helping to take pressure off the main grid, the process would eventually come back around and help to make the data centres themselves less prone to outages.

But there is still some way to go before waste heat utilisation can enter the mainstream. Most data centres still use air-based cooling. Because air isn't an efficient transport medium, consumers of the captured heat need to be located near a centre. Added to this, the infrastructure would need to be upgraded.

"Capturing and reusing heat would require a full overhaul of your entire facility, while [other options] may be less invasive to your hardware set-up," observes Daan Terpstra, executive chairman of the Sustainable Digital Infrastructure Alliance. "But, with a typical hardware refreshment cycle of data centres being somewhere between five and seven years, I think this is an ideal moment to start plotting the chart and placing this at the top of the list."

2

Liquid cooling

With the demand for data rising exponentially, data centres need a lot of energy to stay running – and cool. Specialised computing equipment can emit large amounts of heat. It's important to regulate this to keep the system functioning. Traditionally, that was achieved by creating almost sub-zero, freezer-like conditions, but in recent years the sector has learnt that data centres operate most efficiently at ambient temperatures of between 18ºC and 27ºC. As recently as five years ago, 40% of the total energy consumed by data centres was used in the cooling of equipment. That proportion has since fallen to about 10%.

Although this is the era of air-based cooling, experts agree that liquid cooling – in which heat from equipment is transferred to a liquid and siphoned away – is relatively energy-efficient.

As Lawrence points out: "Air-based cooling pushes the hot air out of the system – you exhaust it. And that is wasted energy."

As opportunities to use waste heat from data centres as an energy source proliferate, liquid cooling is set to become an increasingly important technology. Even in heavily insulated pipes, hot air can't travel very far before cooling too much. Hot liquid, on the other hand, is far more transportable.

"The other thing about direct liquid cooling is it requires very little water," Lawrence says. "It will be easy to use."

3

Artificial intelligence

Several of the world's tech giants have set ambitious renewable-energy targets for their data centres. For instance, Meta, which has more than 20 centres, committed to 100% renewable energy in 2011, followed by Apple, Google and Amazon.

Microsoft has pledged to become carbon-negative by 2030. It has also committed to removing all of the carbon the business has ever emitted, either directly or by electrical consumption, since it was founded in 1975, by 2050. A blog post on its website states: "To reach this, data centres must be part of the solution for broad decarbonisation."

Nonetheless, buying carbon offsets is the method by which many big tech companies are aiming to achieve net zero, which means that they will, in effect, still be using fossil fuels. That situation may change quickly, experts believe, partly because of societal pressure and upcoming legislation.

"Based on the current social and economic climate in continental Europe and the UK, sustainability will become a licence to operate," Terpstra says.

AI is one of the most cost-effective and scalable tools for improving the energy-efficiency of data centres. In 2018, for instance, Google and DeepMind jointly developed an AI-powered recommendation system to control the cooling of data centres, resulting in claimed average energy savings of 30%.

The use of AI can offer more than energy and cost savings. There's also resilience. Alibaba Cloud, for instance, has deployed machine-learning-based temperature alert systems in its global data centre. In July 2021, the firm's principal engineer, Wendy Zhao, told industry publication *Data Centre Dynamics:* "We took hundreds of temperature sensors' monitoring data, using an ensemble graph model to quickly and precisely identify a temperature event due to cooling facility faults. It generated alerts much further in advance and provided the data centre operation team precious time to respond to the fault."

Microsoft is developing an AI system to analyse data and generate alerts to "prevent or mitigate the impact of safety incidents", while Meta is investigating ways in which AI can anticipate how its data centres are likely to operate under "extreme environmental conditions".

4

Microgrids

Most data centres have multiple sources of power so that, if one source fails or goes down, another can keep them functioning. Resilience has always been a primary concern for data centre operators. While the threats and the solutions might be evolving, the ability of a data centre to withstand failures cost-effectively remains paramount.

Microgrids are increasingly being seen as an excellent back-up solution for data centres. A microgrid is an autonomous local energy grid that enables you to generate your electricity, which means that it isn't dependent on the traditional grid. It can not only keep the data centre's power on during grid outages; it can also store electricity and sell it back to the grid.

"So many outages are happening that any critical facility – whether it's a hospital or a data centre – is thinking about how to ensure that it's able to run if the grid goes down, not just for an hour or two but potentially for days or weeks," says Jayesh Goyal, chief revenue officer at Enchanted Rock, a company that's been contracted by Microsoft to develop California's largest microgrid. The facility will use renewable natural gas and provide Microsoft's San Jose data centre with auxiliary power.

What makes microgrids especially noteworthy, Goyal says, is that you can choose how you want to power them. Renewable, natural gas or fuel cells – the choice is yours, constrained only by cost and space. Natural gas is a popular fuel choice for microgrids because of its accessibility and relatively small environmental footprint. But what's exciting to many experts is the opportunity to use hydrogen fuel cells.

In 2020, Microsoft worked with Power Innovations to power an array of data centre servers for 48 hours using fuel cells with a first-of-its-kind hydrogen generator. Hydrogen is described as a clean fuel because water is its sole by-product. But it occurs naturally only in compound form and the cost and technology required in separating it from other elements have been prohibitive. This situation has started to change, though. As it does so, hydrogen-fuelled generators and microgrids start to look like a real possibility.

Terpstra believes that hydrogen fuel cells will need to be used in more than microgrids and back-up generators. Building a data centre fully powered by hydrogen fuel cells is the only route to cost-effectiveness, he argues.

"The calculations I've seen mean that the costs of setting up hydrogen back-ups versus the number of times you'd need them are completely out of balance," Terpstra says. "The run-time on back-ups is too little when compared with the investments required."

Underwater data centres

In 2018, Microsoft ran Project Natick, dropping a data centre containing 855 servers 35m below the sea just off the Orkney Islands. The aim was to insulate the facility from extreme temperature fluctuations and test whether underwater data centres could be reliable and practical while using energy sustainably.

Two years later, the company retrieved its data centre and found that only eight servers were down. Microsoft said that the equivalent figure on land over the same period would have been 64.

Subsea Cloud, which plans to start operating an underwater data centre off the west coast of the US before the end of this year, claims that constructing underwater data centres is cheaper and could reduce carbon emissions by 40%.

In a bid to meet their stated targets, Microsoft and other big companies are experimenting with ways to make data centres more sustainable. While this is to be lauded, many of their experiments are impractical in terms of both cost and scalability, according to Terpstra.

"It may be super-cool to have underwater data centres, but there are so many other solutions possible that would result in the same effect by looking at the reliability and climate impact from a holistic design viewpoint," he says.

For Terpstra and several other experts in this field, it's all about practical measures that can move the needle now – and will continue to create an impact as the infrastructure improves. ●

5