

## Financial Times 번역요약본 ('25. 8/18)

### 1. Donald Trump's 100% tariff threat over chip sector despite relief for Apple : 트럼프의 100% 관세 위협, 애플에 대한 예외에도 반도체 업계 전반에 드리운 그림자 ('25. 8/8)

- 트럼프 대통령이 애플을 향후 부과될 반도체 관세에서 제외하기로 한 결정은, 미국 내 투자를 약속하는 한, 기술 업체와 그 공급업체들이 관세를 피할 수 있다는 신호로 받아들여져 투자자들을 안심시킴. 하지만 업계 관계자들은 이번 관세 조사가 반도체 자체뿐 아니라, 그것을 만드는 장비와 소재, 그리고 반도체를 포함하는 제품까지 포괄하고 있어, 애플처럼 강력한 로비 능력이나 미국 내 공장 건설 자금을 갖추지 못한 많은 기업에는 여전히 큰 타격이 될 수 있다고 경고. 미국 애리조나 주에 1650억 달러 규모의 반도체 공장을 건설 중인 TSMC와 미국 텍사스에 공장을 건설 중인 삼성전자 등은 보호를 받을 가능성이 크지만, 대만의 TSMC를 제외한 중소 파운드리 업체들 (UMC, Vanguard 등)이나 미국의 마이크론이 해외에서 운영하는 공장들은 최대 100% 관세를 피하기 어려울 수 있다는 전망. 결국 기업별, 국가별로 제각각 다른 '누더기식 협의'로 귀결될 가능성이 높다고 보고 있음

### 2. South Korea's Upstage enters global AI race : 한국의 업스테이지, 글로벌 AI 경쟁에 진입 ('25. 8/11)

- 잘 알려지지 않은 한국의 인공지능 스타트업이 미국과 중국의 첨단 시스템에 견줄 만한 성능을 가진 대규모 언어모델 (LLM)을 개발하며, 세계 AI 경쟁에서 뒤쳐진 한국의 추격 의지를 높이고 있음. 서울에 본사를 둔 업스테이지 (Upstage)의 김성훈 대표는 FT와의 인터뷰에서 “예전에는 미국과 중국의 LLM이 모두를 앞서 있었지만 이제는 그렇지 않다”고 말함. 지난 7월 출시된 업스테이지의 솔라 프로

2 (Solar Pro 2)는 독립 벤치마킹 기관인 Artificial Analysis로부터 한국 유일의 선도적 프런티어 모델로 인정받았고, 이 성과로 한국은 미국의 오픈 AI, 중국의 딥시크 외에도 프랑스의 Mistral, 중동의 국영 AI 기업 등 일부 국가만 보유한 첨단 AI 모델 보유국 대열에 합류함. 특히 솔라 프로 2는 매개변수가 300억 개에 불과하면서도 (보통 프런티어 모델은 1-2천억 개의 매개변수 사용) 더 큰 기업들의 모델을 능가함. 한국의 신임 대통령은 한국을 미국, 중국에 이어 세계 3대 AI 강국으로 만들겠다는 목표를 밝히며 700억 달러 이상을 투자하고, 부총리를 겸임하는 AI 장관직을 신설함

### 3. Who pays for the \$3tn AI building boom : 3조 달러 규모의 AI 인프라 건설 붐, 누가 비용을 부담하나? ('25. 8/14)

- Meta는 'Prometheus'와 'Hyperion', 일론 머스크의 xAI는 'Colossus', 오픈 AI는 'Stargate'를 개발 중이며, 각각 천억 달러 이상이 투입되는 세계 최강 슈퍼컴퓨터 건설 프로젝트를 진행하며 차세대 인공지능 시대를 열겠다는 계획임. 그러나 이들 초대형 프로젝트조차 AI 시대를 가동하기 위해 필요한 데이터센터 건설 비용의 일부에 불과하며, 구글, 아마존, MS, Meta 등 빅테크 4대 기업은 올해 3,500억 달러 이상, 내년에는 4천억 달러 이상을 데이터센터 건설에 쏟아부을 예정임. 하지만 자금이 쏟아져 들어오는 와중에도 과잉 공급, 장기 수익성, 에너지 수요 등에 대한 우려가 제기되고 있으며, 한 AI 인프라 금융 담당자는 “많은 데이터센터 건설 시도가 실패할 것”이라며 “지금은 자본시장이 비이성적으로 돈을 쏟아붓는 시기”라고 말함. AI 슈퍼컴퓨터 건설과 데이터센터 확장이 3조 달러 규모의 자본 이동을 촉발할 것이며, 빅테크만으로는 자금 충당이 불가하므로 민간 자본과 부채 조달이 필수적임. 맞춤형 건설 (build-to-suit) 방식이 확산되고 있으며, AI 칩과 냉각 기술의 빠른 진화로 데이터센터와 GPU가 10년 내 고철 덩어리가 될 수 있다는 우려와 함께 제2의 닷컴/통신 버블을 우려하는 목소리도 있음

#### 4. Inside the AI race: can data centers ever truly be green? : AI 경쟁의 이면: 데이터센터는 진정으로 ‘친환경’이 될 수 있을까? (‘25. 8/14)

- 미국, 아일랜드, 독일, 캐나다, UAE 사막 지대 등지에서 AI 어플리케이션과 학습을 위한 전력 수요가 폭증하면서, 화석연료 투자가 되살아나고 있으며, 이는 지구 온난화 주범인 온실가스 감축을 위한 전 세계적 노력과 정면으로 배치됨. 오픈 AI의 Stargate 프로젝트는 텍사스 시설에 360KW규모의 가스발전 용량을 계획 중이고, 머스크의 xAI Colossus 슈퍼컴퓨터에는 35기의 가스 터빈이 설치된 것으로 알려짐. 기술 기업들은 대규모 재생에너지 투자를 내세우며 온실가스 문제를 해결하려 하지만, 상당 부분은 녹색전력 크레딧 (Green Power Credits) 구매에 의존하고 있으며, 이는 실제 소비와 시간, 장소를 일치시키지 않아 온실가스 저감 효과가 제한적이라는 비판을 받고 있음. 국제에너지기구 (IAEA)는 현재 추세라면 데이터센터 전력 사용에서 발생하는 이산화탄소 배출량이 향후 5년 내 두 배로 늘어날 것으로 전망함. 아시아와 중동 역시 데이터센터 성장률이 높지만 대부분 여전히 화석연료에 의존하고 있으며, 배터리 저장, 원자력, 전력망 회계 개혁이 친환경 전환의 핵심이라고 주장함

Trump tariffs

## **Donald Trump's 100% tariff threat looms over chip sector despite relief for Apple**

US probe into semiconductor sector could leave a patchwork of agreements with individual companies and countries

FT reporters

Published AUG 8 2025

---

Donald Trump's decision to exempt Apple from looming semiconductor tariffs has cheered investors who took it as a sign that tech vendors and their suppliers could avoid the levies as long as they pledged US investment.

But industry insiders say the Trump administration's investigation into chips — including the tools and materials used to make them, and products that contain them — will still hit many companies that lack Apple's lobbying muscle or the money to build factories in the US.

“Other vendors of devices from notebooks to smartphones would equally have to prove that their direct investment, as well as that of their supply chain, into the US is material enough for them to get an exemption,” said Nicolas Gaudois, head of Asia-Pacific technology research at UBS. “What about all the vendors who cannot demonstrate [that]?”


He added that all the smaller vendors who cannot demonstrate that would possibly face a 100 per cent tariff.

The so-called [Section 232 investigation](#) is one of nine probes in which the Trump administration is considering sectoral tariffs on industries including cars, steel, pharmaceuticals and drones, on grounds that the US needs to protect its markets for national security.

The semiconductor investigation is one of the most complex and closely watched given chips' importance to the modern global economy, and the deals that the US reaches with some companies — and tariffs it might apply to others — could quickly trigger competitive shifts in the chip industry.

## Taiwan and Malaysia were the largest semiconductor suppliers to the US last year

US imports (\$bn)



FINANCIAL TIMES

Source: US International Trade Commission, Bernstein

On Wednesday Trump said the [chip tariff rate](#) would be “approximately 100 per cent”, adding that it would not apply to those building production capacity in the US.

He named Apple as a beneficiary of that exemption after the iPhone maker raised its pledge for US investments by \$100bn to \$600bn.

Others are not in the same position. An executive at a Japanese chipmaker that supplies Apple said a “huge issue” is that it has hundreds of suppliers “who would have no capacity to go to the US”.

The person said that state-backed agencies that could support Japanese investment in the US were more comfortable with lending to traditional sectors such as steel and shipbuilding. Tokyo has committed to investments worth \$550bn as part of its bilateral tariff deal with Washington.

Executives and analysts said the Trump administration’s opaque, chaotic and transactional approach to tariffs made it hard to gauge what companies have to do in order to be spared.

Taking Apple as an example, one observer called the investments referenced for achieving an exemption for iPhones a “hodgepodge”. They include investment pledges not by Apple itself but by its largest supplier Foxconn, as well as projects that are earmarked for assembly of servers, not smartphones.

“We are having a useful dialogue with [the US Department of] Commerce, but they don’t know what is going to happen,” said an executive at a chipmaker with capacity in the US. “It is unpredictable because Trump decides, he conducts policy extortion, and he is ad hoc transactional.”

Other industry executives involved in discussions with the Commerce department said they expected the probe to result in a patchwork of company- and country-specific deals.

Two said Taiwan Semiconductor Manufacturing Company would be exempt because it is building fabrication plants, or fabs, in Arizona, which the world's largest chip manufacturer has said will eventually add up to investments of \$165bn.

Analysts said Korean smartphone and chip producer Samsung's fab investments in Texas were large enough to give it cover too.

Then there are country-specific arrangements. Sravan Kundojjala, of chip consultancy SemiAnalysis, said Washington was pursuing deals with key allies or manufacturing centres including South Korea, EU, Taiwan and Japan to reduce any disruption to chip supplies.

PLAY | 00:41

Donald Trump threatens 100% tariff on chips © Reuters

The US granted semiconductors, cars and pharmaceuticals imported from the EU the general 15 per cent rate agreed in their bilateral tariff deal, in effect shielding EU products from the upcoming chip tariffs.

South Korea's finance minister told lawmakers this week that the US had agreed a similar provision for chips and pharmaceuticals in their bilateral tariff deal — a commitment in line with US Commerce secretary Howard Lutnick's readout but yet to be confirmed by the White House in writing.


However "precedent suggests any exclusion process established for semiconductors may not be permanent", Kundojjala said. Under a 232 steel industry probe started in Trump's first term, the administration granted country-level exemptions in 2018 but revoked them this year.

Some observers believe the plethora of individual deals could change the dynamics in the sector.

Bernstein analysts said chipmakers other than TSMC could be significantly disadvantaged. For example, fabs owned by US memory chipmaker Micron, as well as Taiwan's smaller chip manufacturers UMC and Vanguard, in Taiwan, Singapore or other countries "may be subject to a 100 per cent rate, unless these countries also negotiate a lower one for Section 232", they wrote in a note on Friday.

## Most semiconductors enter the US inside other things

Semiconductor share of bill materials (BOM) across selected end markets and key products (%)



FINANCIAL TIMES

Source: McKinsey and Company, Bernstein • \*Range of 55-60%

Industry executives are also in the dark over whether the tariffs will be calculated based on the total cost of a device imported into the US or only the part of the bill of materials that reflects the cost of goods sold related to manufacturing the chip outside the US.

Such technical challenges could force the Commerce department to spend more time completing the probe. While Lutnick said in late July the tariffs would come out in two weeks, two people familiar with the probe said the Commerce secretary is now telling companies that several more weeks is needed.

According to Commerce department rules, 232 investigations can take up to 270 days — which would give the department until December.

One of the biggest decisions the administration has to make is whether to apply tariffs to materials and chipmaking tools. Industry experts have warned that by steeply raising the cost of building fabs in the US, this could derail Trump's goal: to bring semiconductor manufacturing back to the US.

Among the more than 150 public comments companies, governments and industry bodies submitted to Commerce on the investigation, warnings about this were one of the dominant concerns.

“The costs that Taiwanese companies face to obtain equipment and components when establishing manufacturing facilities in the US will increase — directly affecting their investment willingness,” the Taiwanese government said in its comments submitted in May.

*Reporting by Kathrin Hille in Taipei, Harry Dempsey in Tokyo, Christian Davies in Seoul and Aime Williams and Demetri Sevastopulo in Washington*

## Artificial intelligence

### South Korea's Upstage enters global AI race

Start-up has produced a large language model that performs as well as advanced systems in the US and China



Upstage chief executive Sung Kim: 'We created a top-tier frontier LLM using only 10% of the number of AI chips that other companies are using, making us much more cost-efficient' © FT montage

**Christian Davies** and **Song Jung-a** in Seoul

Published AUG 11 2025

---

A little-known South Korean artificial intelligence start-up that has produced a large language model that performs as well as advanced systems made in the US and China is seeking to boost the Asian country's goal of catching up lost ground in the global AI race.

Sung Kim, chief executive of Seoul-based Upstage, told the Financial Times that "American and Chinese LLMs used to be way out in front of everyone else, but that is no longer the case".

The claim comes after the company's Solar Pro 2, released in July, became the only Korean LLM to be recognised as a leading-edge "frontier model" by independent benchmarking analysis provider Artificial Analysis.

That advance means South Korea has joined the small number of countries that are home to an advanced AI model. Outside US leaders such as OpenAI and Google and Chinese groups such as DeepSeek, only a handful of groups have world-class frontier models, including France's Mistral and state-backed AI companies in the Middle East.

Kim said South Korea had previously "dropped the ball" on AI due to an excess of caution at its established tech companies, but that Upstage's achievements had put the country back in the race to be the first to develop machines that surpass human intelligence.

According to Artificial Analysis' intelligence index, Solar Pro 2 earned a higher combined score across a range of benchmarks than Anthropic's Claude 3.7 Sonnet Thinking, DeepSeek's V3 and Open AI's GPT-4.1 — putting it just outside the world's top 10 frontier models.

Kim noted that Solar Pro 2's performance had surpassed the models of more established companies despite having just 30bn "parameters", the number of variables used to train an AI system and shape its output.

He said his company had achieved this by deploying a training approach called "Depth-Up Scaling", which involves removing and adding layers of information to a base model to strike a balance between computational efficiency and model performance.

Frontier models typically contain 100bn-200bn parameters, meaning they require more computing power to conduct their calculations. The Grok 4 model made by xAI, which topped the latest Artificial Analysis index, contains 1.7tn parameters.

“We created a top-tier frontier LLM using only 10 per cent of the number of AI chips that other companies are using, making us much more cost-efficient,” said Kim, who served as the founding leader of Korean tech conglomerate Naver’s AI fledgling division before co-founding Upstage in 2020. He added that Upstage intended to develop a new model with 100bn parameters within the next year.

He said Solar Pro 2 specialises in conducting complex calculations for companies in the financial, legal and medical sectors, with its model already being deployed by Intel and the insurance subsidiaries of Korean conglomerates Samsung and Hanwha. Five “big American insurance companies” were also interested in adopting it to assist with underwriting procedures, he added.

Upstage’s emergence offers a shot in the arm to South Korea’s AI ambitions, which until recently rested largely on the supply of hardware such as memory chips and data centre components to leading global players.

The company is also in talks with local Korean chip design start-ups Rebellions and FuriosaAI to deploy their neural processing unit chips for Upstage models to conduct [“inference” tasks](#) — when users request information from LLMs.

“There is a Korean ecosystem developing with a new generation of companies on both the hardware and the software side, and because we all know each other well we can co-operate very closely,” said Kim.

The country’s new leftwing president Lee Jae Myung has declared his ambition for South Korea to emerge as a “top three AI powerhouse” behind the US and China, pledging more than \$70bn in investment and creating a cabinet position of minister for AI, who will also serve as deputy prime minister.

Experts note that South Korea lags behind many other developed countries in data centre investments. It is also suffering from an exodus of AI talent to western countries, according to a report by the Korea Chamber of Commerce and Industry last month.

Kim noted the government had announced a programme to pay up to 85 per cent of annual salaries of up to \$10mn for AI specialists to work in selected leading Korean companies, as part of a bid to encourage local talent to remain in the country.

He added that South Korea also stood to benefit from standing outside the US-China geopolitical rivalry.

By offering to build data centres locally in partner countries, he said, an alliance of South Korean hardware companies and LLM providers backed by the government have a chance to break into markets in countries reluctant to choose between China and the US. “We tell them we can build data centres locally and they love it.”

---

[Copyright](#) The Financial Times Limited 2025. All rights reserved.

---

### **Follow the topics in this article**

---

Technology sector

South Korean business & finance

Artificial intelligence

FT Edit

South Korea

**Big Tech****Who pays for the \$3tn AI building boom?**

With the data centre race in full swing, private capital is joining Big Tech in seeking to capture the rewards

**Tabby Kinder** in San Francisco

Published AN HOUR AGO

---

Meta is building “Prometheus” and “Hyperion”, Elon Musk’s xAI has “Colossus”, and OpenAI is developing “Stargate” — each a more than \$100bn project to build the world’s most powerful supercomputer and usher in a new generation of artificial intelligence.


But each of those gargantuan ventures is just a fraction of the spending required to build the data centres needed to power the [AI era](#): one of the biggest movements of capital in modern history.

“The amount of capital required is absolutely immense,” said Rob Horn, global head of infrastructure and asset-based credit at private equity group Blackstone, which manages an \$85bn data centre platform.

“The scale of the opportunity is exhausting the capital of [any one financial] market, and is requiring an all-of-the-above approach, with private capital playing a very large role.”

Google, Amazon, Microsoft and Meta will spend more than \$400bn on data centres in 2026 — on top of more than \$350bn this year.

Capital expenditure, \$bn



FINANCIAL TIMES

Source: 10-K filings, S&P Global Market Intelligence • Forecasts = shaded area

For years, Big Tech's capital spending grew steadily, focused on cloud, logistics and underlying infrastructure

Then came the launch of ChatGPT in late 2022 and with it a realisation that AI could upend their businesses

Now begins an arms race, where billions are poured into servers, chips and data centres to power generative AI



But as the money floods in, concerns are being raised about overcapacity, long-term profitability and energy demands.

“Lots of people who are trying to build data centres will fail,” said one banker who helps arrange financing for AI infrastructure projects.

“We are in that period where the capital markets are crazy enough to throw money at almost anything. I am curious to see the next phase and whether rationality prevails.”

---


Once seen as a niche part of the real estate market, the frenzied pace of construction has turned data centres into a sought-after asset class.

This year is forecast to break records for development. The US has about 20 gigawatts of operational data centre capacity. Before the end of the year, another 10GW of data centres are projected to break ground globally, and 7GW will reach completion, according to real estate group JLL.

Historically, most of the spend by the “hyperscalers” — Amazon Web Services, Microsoft Azure and Google Cloud — building data centres for their cloud services businesses was self-funded.

But the scale of computing power needed to generative AI is changing that.

## The AI race



This is the third part in a [series](#) exploring the race for AI capacity and the data centres at the heart of billions of dollars in capital investment.

Part 1: [Inside the relentless race for AI capacity](#)

Part 2: [Can data centres ever truly be green?](#)

Part 3: Financing the data centre boom

While internal cash flows largely covered costs of up to \$200bn last year, costs are projected to double this year and increase further next.

Some economists have started to question how much further hyperscalers’ cash reserves can be stretched, and investors want to know when their spending will translate to real revenues from AI services.

Hyperscalers’ generative AI revenues were just \$45bn last year, according to Morgan Stanley analysts — although they predicted revenues would exceed \$1tn by 2028.


This has left a funding chasm that financiers are rushing to fill.

JLL estimates \$170bn of assets will require construction lending or permanent financing this year. Between now and 2029, however, global spending on data centres will hit almost \$3tn,

according to Morgan Stanley analysts. Of that, just \$1.4tn is forecast to come from capital expenditure by Big Tech groups, leaving a mammoth \$1.5tn of financing required from investors and developers.

## Hyperscalers will fund only half of the \$2.9tn of future AI infrastructure

Main financing paths for global data centre spend (2025-2028)



FINANCIAL TIMES

Source: Morgan Stanley Research • \*Asset-backed securities, commercial mortgage-backed securities / \*\* Private equity, venture capital and sovereign wealth funds

The gap will be filled by everything from private equity, venture capital and sovereign wealth to bank loans, publicly listed debt and private credit. But increasingly, the answer is debt.


About \$60bn of loans are going into roughly \$440bn of data centre development projects this year, twice as much debt as in 2024, according to a recent presentation by law firm Norton Rose Fulbright. More than \$25bn of loans were underwritten in the first quarter of this year alone, according to a report by Newmark.

Funding data centres comes not just with the risk that costs overrun, but also that the technology becomes obsolete far quicker than anticipated, requiring new investment that decreases returns for its owner — or forces them to sell at a discount. That means even the deepest pocketed tech groups may want to share the risk, particularly when debt is cheap and readily available.

Deals are being structured in myriad different ways, from structured debt solutions and project finance vehicles to construction loans, asset-backed securitisations and even green bonds to raise money and start building.

## Securitised debt gains foothold in financing AI infrastructure

Estimated annual issuance for data centre securitised credit



FINANCIAL TIMES

Source: Morgan Stanley Research

Meta raised \$29bn — including \$26bn of debt — from private capital investors led by Pimco this month to help fund data centres in Ohio and Louisiana, enabling it to offset high upfront costs and spend its cash on other initiatives with faster returns.

Investors including Apollo, Carlyle, Brookfield and KKR competed in a months-long bidding war to lend to Meta.

Oracle takes a different approach with the 2GW data centre it has signed up to lease in Abilene, Texas. The project is being built by start-up Crusoe and investment group Blue Owl Capital, which have raised about \$5bn of equity from investors and borrowed almost \$10bn from JPMorgan to fund the construction, backed by Oracle's 15-year lease.

In turn, Oracle has agreed to provide OpenAI with 4.5GW of computing power — including from Abilene — in a deal worth about \$30bn a year, which forms the first part of OpenAI's Stargate data centre project in the US. Neither Oracle nor OpenAI will carry the debt raised to build the Abilene site on their balance sheet.

This data centre development model, known as “build-to-suit”, is being replicated by tech companies across the US.



Meta raised \$29bn — including \$26bn of debt — from investors led by Pimco to help fund data centres in Ohio, pictured, and Louisiana, © Meta

“All of the major hyperscalers have self-build programmes. Where third-party developers can add value is when we have sites that are shovel-ready and can deliver on an accelerated timeframe,” said Tim McGuire, capital markets leader for hyperscale data centre developer Rowan Digital.

He said the process of identifying a site, securing the necessary power and building the infrastructure can be a “three-year plus lifecycle. We can cut that cycle in half.”

To get comfortable with the risk involved in a build-to-suit project, lenders, equity investors and developers require hyperscale tenants to sign long-term leases or capacity commitments before they part with their cash. This means they are in effect lending against the creditworthiness of an investment-grade counterparty such as Microsoft or Oracle, a bet that is prompting a race among private capital providers to offer more and larger loans.

In some cases, it also means acquiring data centre developers themselves. Last year, Blackstone bought Australian data centre platform AirTrunk for \$14.9bn, the second-largest data centre deal after KKR and Global Infrastructure Partners’ 2021 purchase of US data centre owner CyrusOne for \$15.5bn. Last week, Apollo struck its own deal, buying a majority stake in data centre builder Stream.

Apollo said data centres would require “several trillion dollars of global investment over the next decade”; it has already deployed \$38bn into data centre-related infrastructure.

But the scale of capital deployment has turned companies with access to land, power or the specialised computer chips used to power AI data centres into potentially very valuable players — if they can prove themselves capable of delivering results for hyperscale tenants.



Oracle has signed a 15-year lease for a 2GW data centre in Abilene, Texas, from which it will partly provide OpenAI with 4.5GW of computing power © OpenAI

“We see multiple developers every week who have undeveloped land but think that they will be signing leases with hyperscale customers tomorrow,” said Sam Southall of Macquarie Capital.

“Essentially everyone with some land and a tenuous path to power is trying to raise capital, but there is a long way to go in order to have credibility with, and be trusted by, these types of tenants.”

Chief among those to have made this play is CoreWeave, a small company founded to mine cryptocurrencies in 2017, a function that required the high-performance Nvidia graphics processing units (GPUs) that have become a key element in training AI models. CoreWeave made the pivot to leasing and operating AI data centres years later, then listed its shares on the Nasdaq exchange in March. It is now worth \$65bn.

The New Jersey-based company funded its shift to AI with large loans, including about \$10bn from Blackstone. Blackstone took security both over CoreWeave’s GPUs — an increasingly popular form of asset-backed financing for AI data centres — and its contracts to lease computing power to Microsoft. Crusoe, which is building Oracle’s Abilene data centre, also started as a crypto-mining company with access to power contracts.


“Data centres are just a fraction of the capital needed,” said Blackstone’s Horn. “If you have a 1GW data centre, it will cost over \$10bn, but all of the equipment costs another \$30bn plus. There is not just a data centre financing opportunity, but an opportunity around equipment, inventory and supply chain finance.”

---

Powering AI makes up less than half of data centre demand at the moment, but it is responsible for almost all of the growth.

The pace of development has drawn comparisons with the telecoms bubble in the late 1990s, when companies laid more than 80mn miles of fibre optic cables across the US in a drastic overestimate of the demand required. The glut meant costs plummeted and many companies failed.

## Rapid expansion in US data centres and power capacity



FINANCIAL TIMES


Source: [DC Byte](#) • Data reflects self-built facilities only. Figures include only live power capacity. NTT data covers all group companies

“People are making forecasts on the assumption that all enterprises will start to use AI technology and pay for it, and pay enough for it to justify the return on investment for all these training facilities,” said a banker who works on data centre deals.

“The conclusion is that we’re all going to be using AI all the time for everything. That’s an incomprehensible world, but one you need to believe in order to not see how this all ends up losing money.”

Big Tech companies stand to lose the most if forecasts about the potential of AI — and the money to be made — are overcooked. By self-funding and owning a large proportion of their data centre capacity, they take on the capital expenditure, operational risks and regulatory burden.

### The AI race



This is the third part in a [series](#) exploring the race for AI capacity and the data centres at the heart of billions of dollars in capital investment.

If demand for AI plateaus, or it emerges that models such as the Chinese start-up DeepSeek’s can be trained far more cheaply, they will be left with huge stranded assets.

Much of the current spending is on data centres capable of training powerful AI models, but as the technology shifts to inference — running those models — the demand for compute will probably drop and assets may become less valuable. Likewise, if power supplies dry up or AI chip supply chains are delayed, returns on investment may suffer.

Part 1: [Inside the relentless race for AI capacity](#)

Part 2: [Can data centres ever truly be green?](#)

Part 3: Financing the data centre boom

“We view cloud services data centre build-outs as fairly robust. We are less confident long-term in the AI training-only locations,” said one executive at a large developer.

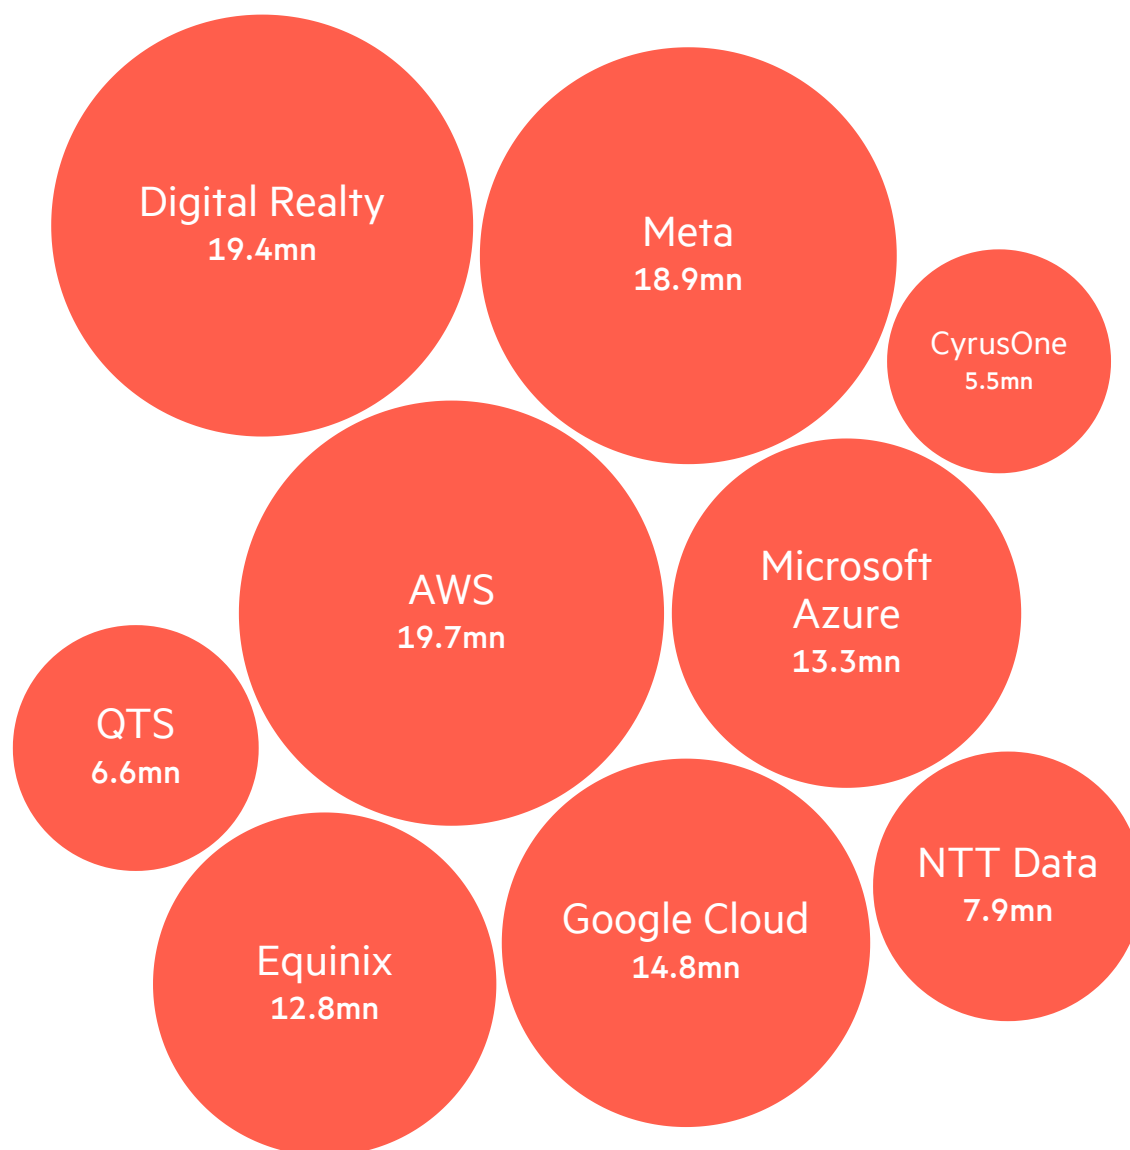
Another banker involved in the sector added: “In five years, you’re not getting a lease renewal at anywhere near the rates you have now. There is a massive overestimation of terminal values.”

The flood of capital into the sector also means increasingly speculative projects are securing funding, including sites that lack an anchor tenant. A larger portion of capital is already being allocated to building data centres for non-investment grade tenants, for example CoreWeave and OpenAI, as well as smaller AI start-ups.

---

## The size of AI infrastructure is immense... and growing

Square footage, Q4 2024



FINANCIAL TIMES

Source: [DC Byte](#) • Data reflects self-built facilities and live square footage only. NTT data includes all group companies

Some landlords, such as Blackstone and infrastructure investor DigitalBridge, are increasingly seeking to extract cash from assets once they are fully built with contracts in place by turning to securitisation deals.

But some investors already view that debt as too risky. One large buyer of securitised debts told the Financial Times they were avoiding such deals because of concerns the properties would be obsolete by the time the debt matured, leaving lenders with assets of questionable value.

“People are lending at high loan-to-value ratios and there is obsolescence risk,” the person said.

Loans secured against GPUs are also at the riskier end of the data centre development financing market and are becoming more common after the success of CoreWeave, with more lenders financing former crypto companies that possess the chips. But Nvidia, the largest manufacturer and designer of AI chips, frequently produces faster and higher-performance chips that risk making its older generations obsolete and devaluing lenders’ security.

Obsolescence is also a risk for the data centres themselves. Many are being built to house Nvidia’s latest Blackwell chips, which require complex liquid cooling systems. The technology is changing so fast that it is feasible future chips will require newer cooling methods.

“There’s a risk in 10 years you just have a shed with obsolete GPUs and cooling infrastructure that is unfit for purpose and you may as well start again,” one person in the data centre sector said.

Still, the hyperscalers can afford the risk. They have the economies of scale to protect against at least some losses. The position of those who borrowed large sums to fund the build-out is more perilous.

“Microsoft and Amazon don’t know what they need. They’re just gobbling up everything because popular opinion is that it’s a winner-takes-all market,” said an executive at a data centre leasing group. “The people who are going to suffer from a pullback . . . are the data centre companies who are overleveraged.”

*Additional reporting by [Rafe Uddin](#) and [Antoine Gara](#). Data visualisation by [Ray Douglas](#)*

Artificial intelligence

**Inside the AI race: can data centres ever truly be green?**

Energy demand for training machines and running apps is driving a surge of investment into fossil fuels

FT reporters

Published AUG 7 2025

---

**In the hills of what used to be coal country in western Pennsylvania,** artificial intelligence is giving new life to a decommissioned coal plant.

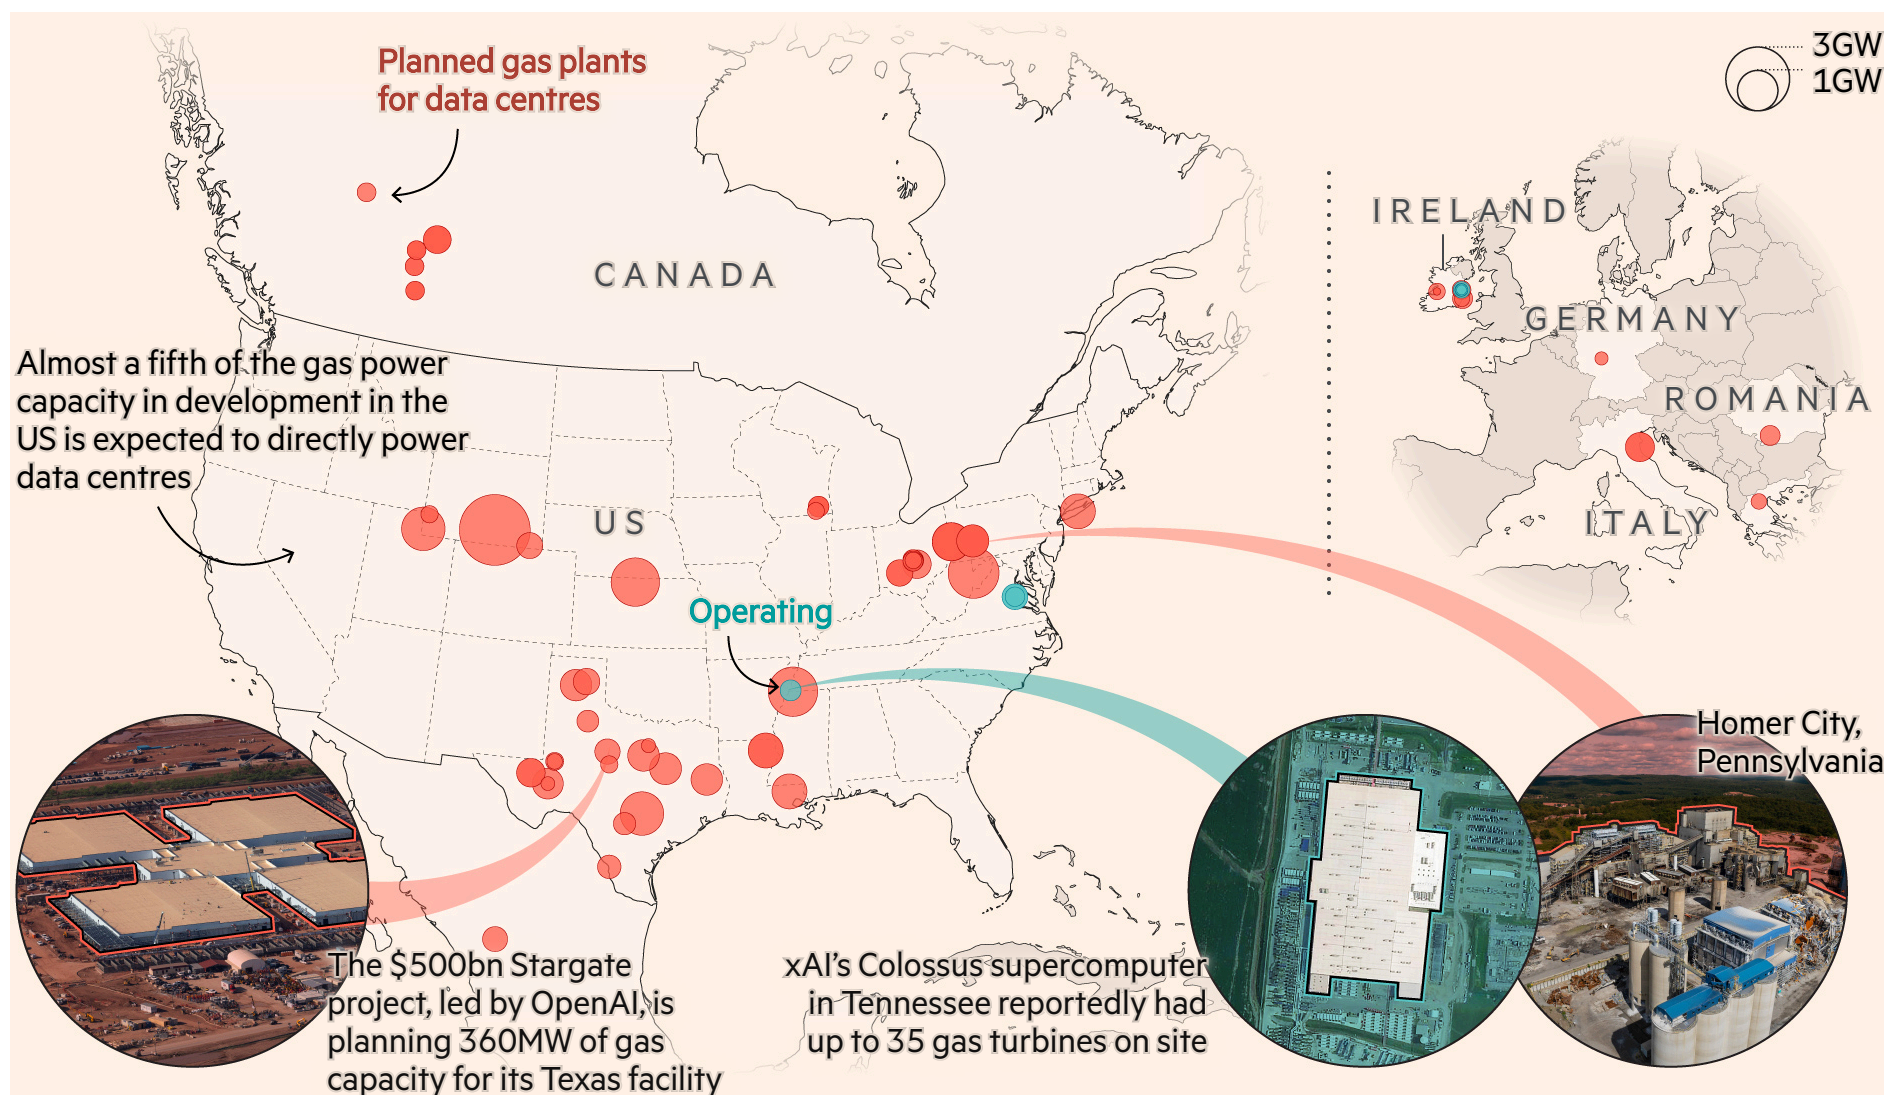
Now being converted into a giant gas-powered facility, it will supply new US data centres needing “round-the-clock reliable power”, said the company behind the project, Homer City Redevelopment. “This project is an investment in the future.”

The plant is expected to be the largest in the US, producing up to 4.5 gigawatts of electricity — enough to power several San Francisco-sized cities.

RESUME

Drone footage of the decommissioned coal plant in Homer City that is being redeveloped into a gas-fired power plant, supplying electricity to data centres © Jae Minard / FT

This is just one of more than 85 gas-fired power facilities in development around the world to supply data centres and meet their burgeoning energy demands from AI, Global Energy Monitor research shows.



These captive gas plants are power units built to supply electricity to a specific facility – in this case data centres – independently of the public grid. Typically such plants are located on-site or near the facility being powered  
Source: Global Energy Monitor

From the deserts of the United Arab Emirates to the outskirts of Ireland's capital, the energy demands of AI applications and training running through these centres are driving the surge of investment into fossil fuels.

This threatens to subvert global efforts to cut greenhouse gas emissions, which are the main driver of [climate change](#). To limit rising temperatures, global emissions must fall by about half by 2030.

Tech companies point to their record as big investors in clean energy in an effort to deal with their emissions, but much of this is backed by green power credits.

This is a controversial form of investing in the build-out of clean energy because it is far from guaranteed to compensate for actual energy consumption.



Data centres' climate problem goes well beyond new gas infrastructure. Just as relevant is the added pressure on existing electricity grids still in large part powered by coal, oil and gas.

Fossil fuels will supply more than 50% of data centre power until 2030

...driving up carbon dioxide emissions this decade

Data centre electricity generation, TWh

Data centre emissions, mn tonnes CO<sub>2</sub>



Source: IEA

In the UAE, the petrostate that is positioning itself as an AI hub, a flagship 5GW data centre cluster project involving OpenAI's Stargate is expected to use the equivalent of one billion cubic feet of gas a day.


The energy demands of AI's learning phase are high and volatile, while those of a data centre are high and constant. Both are poorly matched with the patterns of solar and wind, which wax and wane during the course of a day and a year, sometimes unpredictably.

Mike Hemsley, deputy director at the Energy Transitions Commission think-tank, said: "AI needs a spiky profile while doing deep learning," which means powering a data centre using clean energy is still a purely "theoretical" idea. Data centres are "scrambling for gas turbines because that's the most obvious choice", he added.

Even as tech companies have become some of the world's biggest investors in clean energy, their investments have not solved the conundrum of how to ensure stable supply during cloudy or windless times of the day.

"Our data centres have to run 100 per cent of the time so we can't even consider using [renewables](#) as a primary source," Simon Tusha, founder of US data centre provider TECfusions, told the Financial Times.

## AI data centres



This is the second part in a [series](#) exploring the race for AI capacity and the data centres at the heart of billions of dollars in capital investment.

Part 1: [The relentless race for AI capacity](#)

Part 2: Can data centres ever be green?

Part 3: Financing the data centre boom

Based on current deployment trends, the International Energy Agency expects greenhouse gases from data centre energy consumption to double in the next five years as a proportion of global emissions from burning fossil fuels. The rush for gas is also leading to a global supply crunch for new turbines.

“While the consciousness around the environment comes in and comes out, it will come back in,” said Eamon Ryan, Ireland’s former climate minister. “I don’t think the future is going to be ‘burn, baby, burn’, gas-fired data centres everywhere. That’s not viable because that kills our planet — and explain that to your 20-year-old.”

---

### **Ireland, Europe’s capital for “hyperscaler”**

**cloud data providers**, in May said it was [not on track to hit its climate targets](#) because its emissions are falling too slowly across electricity, transport, industry and agriculture.

Data centres account for more than a fifth of the country’s electricity demands, in a grid system where more than half of power production comes from fossil fuels.

The growth of these centres in the previous decade was concentrated around Dublin, leading to a de facto ban on new facilities since 2021 because of grid capacity issues.


Despite this, the Big Tech companies including Amazon, Meta and Google, which operate some of Ireland’s power-hungry centres, report their energy use is already, or is on track to be, “100 per cent matched” with renewable energy at a global scale this year.

To bridge the gap between the use and supply, the tech companies buy clean energy credits, or can make these investments directly through power purchase agreements.


For example, when a company pays for wind or solar energy to be produced for a decade, it earns the right to claim credits representing the additional clean energy.

But these arrangements are usually purely financial in nature — meaning the data centre is still supplied with power by the central electricity grid.


Credits do not have to match the location where energy is being used, nor the time of day when it is needed.




**Winter, Jan 23**



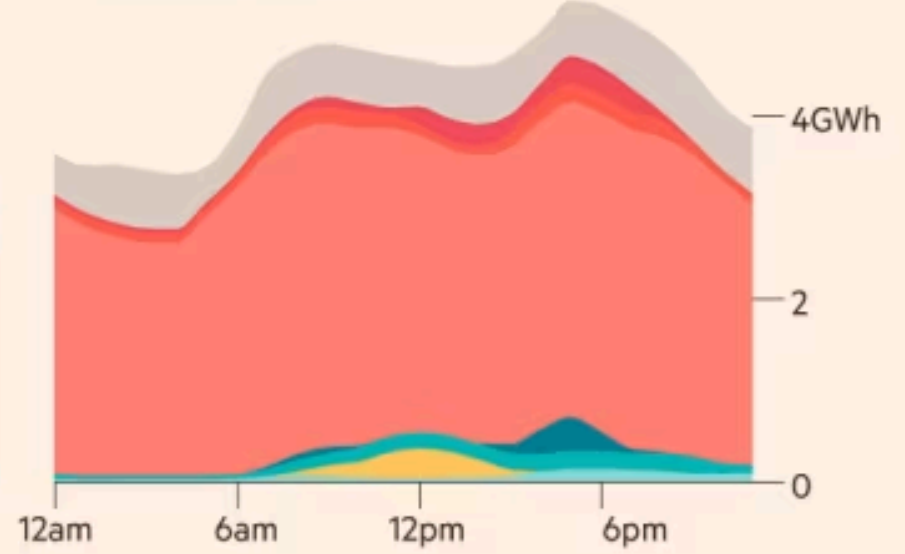
**Spring, Apr 18**



**Summer, Jul 29**



**Autumn, Nov 12**



Source: EirGrid and SEMO data via Green Collective

Big Tech companies can claim their power is fully matched by renewables. In fact, their data centres remain powered in large part by the local grid system.

Take Ireland. In 2024, gas was the largest source of electricity generation.

And on some days,  
almost all of its power  
came from fossil fuels.

On a particularly windy day in January, more than three-quarters of the country's power came from renewables.

But wind is not always consistent, which means gas use varies across days and seasons.

For a data centre in Dublin, for example, the credits claimed for its clean energy use can come from any generation at any time anywhere in Europe.

This strange dynamic means clean energy investments by companies operating in Ireland do not necessarily clean up the country's grid.


The energy investments were “not decarbonising our electricity supply, just serving Big Tech”, said Hannah Daly, professor in sustainable energy at University College Cork.

The most polluting forms of energy are typically used to serve peak data centre demand, including rapid but inefficient oil boilers and open cycle gas turbines, Daly added. This happens regardless of investments in wind and solar.

---

## The surge in electricity used by data centres in Ireland has outpaced growth in renewables generation over the past five years

Change vs 2019, terawatt hours



Source: Ireland Central Statistics Office; Ember

Microsoft blames the failure of grid systems to keep up “even where we have successfully decarbonised our footprint under market-based principles”.

Amazon Web Services also backs the system of buying certificates representing green energy.

It gets some of these credits by investing in new clean energy locally, including in [Ireland's](#) first wind farm to be built without public subsidies, which came online in 2021. But it can also offset its Irish energy use with credits from investments elsewhere in Europe such as Greece or Poland.

This mismatch is replicated across many parts of the world, where tech companies are claiming to build “greener” data centres, while drawing energy from grids that rely on fossil fuels more than on clean energy.

Across south-east Asia, from the Riau Islands of northern Indonesia, to the southern Malaysian state of Johor and neighbouring Singapore, new data centres are springing up alongside vast solar farms.

On aggregate, however, across south-east Asia the data centres are mostly still reliant on fossil fuels, including Singapore's, which was 92 per cent powered by natural gas in 2023.

Data centres will account for up to 30 per cent of power demand in Malaysia, 14 per cent in the Philippines and 12 per cent in Singapore by the end of the decade, according to energy think-tank Ember.

A parallel challenge is that emissions from making the chips needed to operate AI models grew more than fourfold last year, the campaign group Greenpeace estimated. The electricity used came mostly from fossil fuels in east Asia's power grids including in South Korea and Taiwan.

"You have a catalyst for growth in data centres that will outpace the renewable energy market," said Jon Healy, Europe, Middle East and Africa managing director for data centre service provider Salute. "There is a mismatch in that growth."

---

**Even if the rollout of wind and solar power** eventually equalled the pace of growth in data centre energy demand, data centre operators argued there would still be a problem: reliability.

Irish group Echelon has since 2021 been developing a data centre campus beside the motorway that encircles most of Dublin, for a tech client that it declined to name. Solar panels are neatly lined up on the centre's rooftops. It will also buy clean energy credits.

But atop an energy centre nearby, there are four pairs of twisted metal stacks. These gas turbines will stand ready to provide power for times when electricity from the national grid is not available.



Gas turbine exhaust stacks on the campus being developed by Echelon, one of Europe's largest data centre operators © Paulo Nunes dos Santos/FT

Its head of energy systems, Cormac Nevin, said power networks needed to be supported by "fossil fuels or some conventional generation".

"We do need to be cognisant of the fact that we can't get rid of fossil fuels in the morning."

The rest of the time the data centre will plug in directly to Ireland's grid, contributing to demand for gas-fired power.



Cormac Nevin, head of energy systems at Echelon, beside solar panels on the data centre rooftop for power on cloudless days © Paulo Nunes dos Santos/FT

Nuclear is another way of making sure power is reliable, or “firm”, around the clock. Google, Amazon, Meta and Microsoft have all struck deals recently with [nuclear power plant operators](#). But outside China, it comes with local community concerns, delays in obtaining permits and high upfront costs.

Storing clean power in large batteries will be a significant piece of the energy puzzle, experts say, especially as the batteries become cheaper and more ubiquitous.

In sunny areas such as Los Angeles, or Muscat in Oman, the cost of combining solar energy and battery storage to get to 97 per cent clean energy capacity has fallen by more than a fifth in the past year, Ember reports.

However, currently the technology will only go so far. Each additional percentage of reliable clean energy requires disproportionately more spending on battery storage, leading experts to conclude tech companies must be more flexible in when and where they demand energy from grids.

Some hyperscalers are experimenting with “load-shifting” as far as consumer demand and laws will allow, navigating around sensitive commercial or private data storage issues.

Google changes the time it performs some non-urgent computing tasks, or switches the tasks to other jurisdictions, to make the most of times of day and places where clean energy is flowing.

But this kind of load-shifting can be complex and expensive.

“My concern as a user is what am I going to get when the wind don’t blow and the sun ain’t shining . . . because I can’t switch [the data centre] off,” said Pip Squire, head of energy and sustainability at Ark data centres, at a conference this year. “Actually shifting load is a real pain in the ass.”


Sarah Smith, a research scientist at the Lawrence Berkeley National Laboratory, said while it was “perfectly possible” for AI model builders to “throttle or pause a training run” to avoid peak energy demand periods, they are not incentivised to do so by governments.

Some large energy users and climate policy experts argue the solution ultimately lies in the complete reform of the dry but fundamental issue of energy accounting.

Accounting, which better reflects the physical reality of electricity grids, would push data centre operators to invest in renewable energy from local sources that match the timing of local energy usage.

## Accounting techniques continue to distract from true surge in tech company emissions

Million metric tonnes of carbon dioxide equivalent, from company reports



Tech companies work out the 'location-based' carbon footprint of the energy used to power their operations and data centres, based on the mix of fossil fuels and renewables in each local grid. Then, to arrive at a 'market-based' figure, they adjust this number downwards using the lower emissions linked to their clean energy investments in their own or in other grids. Proponents of market-based accounting argue that this incentivises investments into hard-to-decarbonise grids, and more accurately reflects the carbon footprint of a company's energy investments

Source: Company sustainability and auditor reports

“The tech sector has a climate strategy crisis,” said Thomas Day, a policy specialist at the NewClimate Institute non-profit organisation.

A move towards more “credible and useful” reporting that reflects energy grids’ hourly energy flows on a round the clock basis would at least bring transparency, he said. “Anyone who claims 100 per cent renewable is pulling your leg, so we’d rather have a system that reflects reality.”

A handful of other large energy consumers including the pharmaceutical company AstraZeneca and Google have said they back the idea of matching energy investments to energy use on an hourly basis.

This gives them an incentive to support grid resilience.

Google in July said it had invested in CO<sub>2</sub>-based battery storage that can dispatch clean energy for up to 24 hours compared with the typical four hours of lithium-ion batteries.

Some companies such as Microsoft have said they will only buy credits “bundled” in with their own power purchase agreements. These are generally more expensive and more directly tied to “additional” clean power than those traded on open markets.

AWS defended a system that rewards investments in clean energy through certificates as “the fastest, most cost-effective and scalable way to leverage corporate carbon-free energy”, but said the way certificates were accounted for should be reformed.

Even the strongest critics of tech companies are hopeful that after the present gas-fuelled boom, investments in battery storage and demand management could help the climate-friendly claims come closer to reality.

“We can and will and should work with data centres. It’s not a fatwa against data centres,” said Ryan, the former Ireland climate minister. “The industry that learns how to do it in a climate responsible way will have a comparative advantage.”

*Reporting team: [Kenza Bryan](#), [Jana Tauschinski](#) and [Nassos Stylianou](#) in London, [Eva Xiao](#) in Homer City, [Jude Webber](#) in Dublin, [Chloe Cornish](#) in Dubai and [Owen Walker](#) in Singapore*

*Additional reporting by [Rafe Uddin](#) in San Francisco*

## Climate Capital

---



Where climate change meets business, markets and politics. [Explore the FT’s coverage here.](#)

Are you curious about the FT’s environmental sustainability commitments? [Find out more about our science-based targets here](#)