

The Digital Sovereignty Brief #3 — Off, back on, out of our hands

The US switched Anthropic's top models off worldwide, then back on eighteen days later; told OpenAI to vet its new release through Washington; and Britain had no say in any of it.



MARTIN DE SAULLES

JUL 02, 2026

THE DIGITAL SOVEREIGNTY BRIEF • ISSUE 3

Off, back on, out of our hands

A US order switched Claude off worldwide, then back on eighteen days later — and set the terms of access for everyone else.

18 days

Claude switched off worldwide, then back on

~24

OpenAI partners vetted by Washington

30 days

US pre-release look at every frontier model

16 GW

Global data-centre supply — record scarcity

Sources: US Dept of Commerce / BIS; Federal Register (EO 14409); CBRE, Global Data Center Trends 2026.

© Information Matters 2026

The week in six lines:

- **A US order switched Claude off worldwide — then switched it back on.** The Commerce Department made Anthropic disable its two most capable models on 12 June and lifted the order on 30 June; the access of every UK organisation running on Claude was toggled by a foreign government, with no notice and no British involvement in either decision.

- **OpenAI's best model now ships through a US government gate.** GPT-5.6 is going first to a couple of dozen partners that Washington vets one by one, before any wider release.
- **Both run on the same machinery.** A US executive order signed on 2 June lets the federal government take up to thirty days' access to the most capable models — and help pick who gets them next — before wider release.
- **You cannot build your way out of this quickly.** CBRE reports global data-centre supply at record scarcity; London, Britain's only first-tier hub, is the tightest it has been and constrained by power, not money.
- **Britain's compute answer is brokerage, not independence.** DSIT has gone to market for a single broker to aggregate existing clouds — most of them American — into the national AI compute resource.
- **The thread:** the systems, the gate that releases them, and the ground they run on are all controlled from outside Britain.

The through-line: every control point over the frontier AI that UK organisations increasingly depend on — the models themselves, the gate that decides who gets them, and the physical capacity underneath — sits outside Britain, and nothing that happened this fortnight was decided here.

Thanks for reading The Digital Sovereignty
Brief! Subscribe for free to receive new
posts.

✓ **Subscribed**

Ownership & control

The US government switched Anthropic's most capable models off for the whole world, then switched them back on eighteen days later.

On 12 June the Commerce Department's Bureau of Industry and Security ordered Anthropic to suspend its two most powerful models, Fable 5 and Mythos 5, for any foreign national. Because Anthropic cannot separate US persons from everyone else in real time across a user base in the hundreds of millions, the practical result was a hard global shutoff. The trigger was a jailbreak found by Amazon — Anthropic's largest investor and one of its trusted testing partners — in which the model was asked to read a codebase and point out its security flaws. On 30 June, after Anthropic shipped a new safeguard aimed at that specific technique, Commerce withdrew the order — having already cleared Mythos 5 for a set of approved US organisations on 26 June — and Anthropic began restoring Fable 5 to users worldwide on 1 July. Anthropic's position throughout was that a single narrow jailbreak should not justify recalling a model used by hundreds of millions, and that applying such a standard across the industry “would essentially halt all new model deployments.”

So what? this is the clearest illustration of cut-off risk we have had. A UK hospital, law firm or bank standing on Claude lost the service on 12 June for reasons decided entirely in Washington, had no route of appeal in Britain, and got it back on 1 July for reasons that ran through the US Commerce Department and a US cloud provider's security team. No data moved and the contract did not change hands — but the ability to keep operating was turned off and on by a foreign government. The reversal is not the reassuring part of this story. It is the confirmation that the switch exists, works, and can be used again. [Anthropic statement \(official\)](#) · [CNN Business](#) · [CNBC](#)

OpenAI will hand its newest model first to partners the US government has approved. In late June OpenAI began a staggered release of its GPT-5.6 family, making it available at first only to a small

circle of “trusted partners” — each of them approved by the federal government — before any broader rollout. The company has said the requests to limit distribution came from the Treasury, the Commerce Department and two White House offices, on concerns about cyber capability and leakage to adversaries. OpenAI’s chief executive told staff the arrangement is “not our preferred long term model.”

So what? read alongside Anthropic, this is the same story from the other end. The two frontier-model providers UK organisations lean on most are now releasing their best systems through a US government vetting gate, one customer at a time. A British buyer’s access to the newest capabilities now depends on whether Washington approves the counterparty — and there is no equivalent UK process, and no British seat at the table where the approved list is drawn up. [Forbes](#) · [SiliconANGLE](#)

Data & infrastructure

The compute any home-grown alternative would need is in the shortest supply on record. CBRE’s Global Data Center Trends 2026 reports that supply across the sixteen largest markets reached sixteen gigawatts in the first quarter, up a quarter on the year — and that average vacancy still fell, to 6.7% from 8.3%, as demand outran the build-out. London remains the largest data-centre market in Europe, but CBRE names it, alongside Frankfurt, among the hubs where power availability, grid constraints and local opposition are actively limiting new supply.

So what? when the answer to model dependence is “build UK and European capacity,” these numbers show why that answer is slow. Britain’s only first-tier hub is tighter than it has ever been and held back by power and planning rather than by capital, so a sovereign-cloud ambition runs into a grid-and-land bottleneck that no announcement can clear inside a procurement cycle. The physical layer sets the pace,

and right now it is set slow. [CBRE — Global Data Center Trends 2026](#) · [CBRE press release](#)

Britain's national compute plan is to broker other people's clouds.

DSIT's live procurement — "AIRR Expansion AI Cloud Compute," worth about £214m before VAT (£250m with) — seeks a single managed-service provider to act as a strategic broker, aggregating capacity from several cloud providers into the AI Research Resource, the state-backed compute vehicle at the centre of the UK Compute Roadmap. The contract runs from this summer to 2029.

So what? the design choice is the story. The UK's headline compute vehicle is being assembled as a brokerage layer over existing clouds, most of them the same US providers the week's other news shows can be reached from Washington. That is a reasonable near-term move given the capacity reality above, and it should improve the terms Britain buys on — but it is aggregation and access, not control, and it should not be described as sovereign compute in the sense that matters here. [Find a Tender — AIRR Expansion AI Cloud Compute](#) · [DSIT — UK Compute Roadmap](#) · [UKAuthority](#)

Reports & analysis

Independent analysts read the Anthropic episode as a precedent, not a one-off. Three credible bodies published on the shutoff this fortnight. The Center for Strategic and International Studies asked what comes next now that Commerce has shown it will restrict access to a deployed commercial model; Lawfare described the episode plainly as a "kill switch for frontier AI"; and Tech Policy Press argued the block set an export-control precedent for the whole field. Their shared point is that the instrument was not any new AI-specific law but the existing export-control apparatus, turned on a service millions already use.

So what? the 30 June reversal does not retract the precedent — it demonstrates the mechanism and leaves it in place. For any UK

organisation standardising on a US frontier model, the analysts' conclusion is the one our framework keeps reaching: continuity of access is a function of US executive discretion, and Britain has no exemption, no notification right and no way to contest it. Treat these as argument rather than evidence, but the argument is well made. [CSIS](#) · [Lawfare](#) · [Tech Policy Press](#)

Regulation & policy

One US executive order sits behind both stories. Executive Order 14409, signed on 2 June, directs agencies to design — within 60 days, so by around 1 August — a “voluntary framework” (its words) with two moving parts. First, developers give the federal government access to a “covered frontier model” for up to 30 days *before they release it to other trusted partners*. Second, the government “collaborate[s]” with the developer to *select those trusted partners*. Which models count as “covered” is decided by the Director of the National Security Agency through a classified benchmarking process. The order states it creates no mandatory licensing, pre-clearance or permitting requirement.

So what? this is the machinery this fortnight ran on. A US framework now shapes the release cadence and access terms of the frontier models UK public bodies and firms depend on — and the Anthropic and OpenAI events show what “voluntary” means when the same government holds export-control powers in its other hand. Britain has no equivalent instrument, no reciprocal access, and no standing in the NSA process that decides which models are gated. It is worth reading as the wiring diagram for everything above. [Federal Register — EO 14409](#) · [The American Presidency Project](#)

Britain's one live lever over the cloud giants now has a deadline. The regime that lets the Treasury designate the largest technology suppliers as “critical third parties” to UK finance — bringing them under Bank of England, PRA and FCA oversight — still has no names on it. This week, though, the position moved: the Treasury Select Committee has

recommended designating the major cloud and AI providers by the end of 2026, and the Treasury says it is gathering evidence and expects to make its first designations this year.

So what? the hyperscalers that UK banks and insurers cannot operate without still sit outside direct supervision, but for the first time a decision is expected inside six months. A designation would be the first real UK instrument aimed at concentration in the financial system's technology base — and the moment it lands is the moment our Microsoft 365 assessment and the wider cloud-concentration picture come up for review. [Bank of England — Critical Third Parties](#)

From Information Matters this week. Four pieces in our Applications & Infrastructure series went live on 30 June, each mapping who controls one layer of the UK stack:

- **Cloud & data centres** — why a London data centre run by a US company is still within reach of US law, and the roughly 84% of UK business software that runs on the three American clouds. [Read.](#)
- **Identity & access** — the log-in layer that gates everything else, twenty-one products deep, and why it is the fastest switch-off risk in the stack. [Read.](#)
- **Payment rails** — Britain owns the bank-to-bank rails but rents the cards: Visa and Mastercard carry almost every UK card payment, and a Mastercard subsidiary runs the UK's own Faster Payments platform. [Read.](#)
- **AI models** — whose engine reads your data, and the Microsoft 365 Copilot disclosure that its Anthropic subprocessor sits outside the EU Data Boundary. It is the direct backdrop to this week's lead. [Read.](#)

Coming up at Information Matters.

- **Fourteen Sector Exposure Profiles — publishing next week.** The framework behind this week's four Applications & Infrastructure pieces, applied sector by sector and scored on our exposure scale — who controls the stack, whose law can reach the data, and how hard it would be to switch. Ordered most-exposed first: **Construction & property** (3.7), **Agriculture** (3.7), **Insurance** (3.6), **Manufacturing** (3.6), **Media & creative** (3.6), **Local councils** (3.6), **Retail & e-commerce** (3.6), **Accountancy** (3.5), **Wealth & asset management** (3.5), **Hospitality & leisure** (3.5), **Logistics & transport** (3.5), **Charity & non-profit** (3.5), **Education** (3.4) and **Energy & utilities** (3.4). [The Sector Exposure hub](#).

Thanks for reading The Digital Sovereignty
Brief! Subscribe for free to receive new
posts.

✓ **Subscribed**

*The Digital Sovereignty Brief · Information Matters · © Information Matters
2026*