

Data Centre Briefing

April 04, 2026

Global

Key themes:

French state buys Bull from Atos for €404m; AirTrunk targets 2026 Singapore REIT IPO raising \$1.5B; DOE rescinds ALARA as Amazon cites \$1B nuclear investments; Pennsylvania data centre backlash: O'Hara Township summit, 93dB and 5m gpd

France just drew a thick line around “sovereign compute.” In a move that would’ve sounded unthinkable a few years ago, the state has bought supercomputer firm Bull from Atos for **€404m**, explicitly citing national security and nuclear weapons research ties at **CEA-DAM**. If you’ve been treating AI and HPC as “just another capacity cycle,” today’s reminder is blunt: governments now see this stack as strategic infrastructure.

The Big Stories

The French government has [nationalized supercomputer firm Bull from Atos for €404 million](#), arguing the acquisition is needed to protect national security because Bull’s systems are used by the CEA-DAM nuclear weapons research lab. The story also flags Bull’s role in major exascale efforts, including building the **1-exaflop Jupiter system** and participating in the consortium delivering the **Alice Recoque exascale supercomputer in 2027**. The signal here is bigger than one carve-out: European states are increasingly willing to use ownership—not just regulation—to secure critical compute and the supply chain around it.

AirTrunk is reportedly lining up capital-market optionality in a very data-centre-native way: [a Singapore REIT IPO in 2026 to raise \\$1.5B](#). The context matters: Blackstone and CPP Investments agreed to acquire AirTrunk for **\$16.1bn** in

2024, and the company is described as having **>800MW committed capacity** plus land that could support **>1GW** of future growth. Read this as a play to recycle capital while keeping growth velocity—especially useful if debt markets wobble or capex keeps inflating.

US nuclear policy just took a turn that will land uncomfortably in the same conversations as hyperscaler power strategy. The Department of Energy [rescinded the ALARA radiation-exposure directive](#) via a January 9 memo to speed nuclear development under President Trump’s executive order, prompting criticism about reduced worker protections. The connective tissue to data centres is explicit: hyperscalers including **Amazon, Meta, and Google** are backing **small modular reactors**, and Amazon says it invested **more than \$1bn in nuclear projects last year**. If nuclear is becoming part of the “credible power” toolkit for AI campuses, policy shortcuts (and the backlash they trigger) become a non-trivial execution risk.

AI is also changing how the industry defines “reliability,” not just how much power it needs. An industry analysis argues for a shift to [workload-specific “precision resilience” designs](#), separating **energy-optimized training campuses** from **high-availability inference sites**, and pushing standard reference designs, offsite manufacturing, and modular upgrade paths. The practical takeaway: the old default of designing everything like a bank-grade facility is colliding with AI’s economics—especially where power efficiency and speed to deploy matter more than traditional uptime dogma.

Community pushback in the US is becoming more organised and more technical. [Community Action Works is hosting a day-long summit in O’Hara Township on April 18](#) to help residents respond to dozens of proposed Pennsylvania data centres, citing concerns like diesel generator pollution, noise reportedly measured up to **93 decibels**, and water use up to **5 million gallons per day**. This isn’t generic NIMBY noise; it’s a sign that permitting fights are gaining playbooks—workshops on zoning, mapping proposals, and campaign organising—which can shift timelines and raise soft costs even where land and power look available on paper.

Behind the Headlines

Greenpeace is trying to reframe AI buildout as a legitimacy problem, not just a grid problem. In a new briefing, [it warns of AI's environmental and "democratic" costs](#), citing claims that AI data-centre electricity demand could be **11× higher in 2030 than 2023**, and arguing that **74% of industry climate-benefit claims are unproven**. The document stitches together disparate pushback examples (from local councils to legal challenges) to make a single point: community consent and trust are becoming constraints. For investors, the key isn't whether you agree with the framing—it's that this narrative is getting packaged for mainstream political use.

The insurance market is quietly adapting to the way data centres now combine physical, cyber, and business interruption exposures into one messy risk profile. A piece on [multi-line versus monoline coverage](#) notes insurers offering bundled policies, with examples of expanded capacity and offerings in January 2026 from **Aon (~\$2.5bn)**, **FM (~\$5bn)**, and **ATA (up to \$750m per facility)**. Bundling can cut admin burden and reduce coverage gaps, but it can also cap payouts or limit customization versus specialist monoline coverage. In plain terms: as the asset class scales, insurance is becoming another place where "standardisation" can hide nasty edge cases.

AWS's latest quantum-adjacent demo is really a story about classical compute staying in the loop longer than people assume. [AWS and partners simulated a 97-physical-qubit error-correction code on EC2](#), completing a full syndrome-extraction cycle in about **one hour** on a single **Hpc7a** instance with **96 vCPUs**, using a hardware-calibrated digital twin. The point isn't that cloud makes quantum "easy"; it's that progress on fault tolerance still leans heavily on HPC-scale simulation. That keeps cloud and data-centre infrastructure central to the quantum roadmap—even before meaningful quantum workloads show up in production.