

Data Center Briefing

March 24, 2026

Global

Key themes:

OpenAI talks Helion fusion: 5GW by 2030; Washington SB 6355 transmission authority, SB 5982 data-center generators; Ireland CRU decision challenged: 80% renewables, six-year window; Dell PowerCool eRDHx claims 74% cooling energy reduction

OpenAI is reportedly negotiating to buy fusion electricity from Helion — talk of **up to 5GW by 2030** and **50GW by 2035** is the kind of number that stops being “energy procurement” and starts looking like industrial policy-by-contract. If even a fraction of that is real, it’s a tell that the biggest AI players are already planning around a world where today’s grid queues and PPAs simply can’t keep up. And it lands on a day when US states are tightening the rulebook on how data centers plug in — and how they explain their power and water use.

The Big Stories

[OpenAI reportedly in talks to buy fusion power from Helion](#) is the headline grabber: sources say the deal could guarantee **up to 5GW by 2030** and **50GW by 2035**. Helion has a Polaris prototype that has reportedly demonstrated DT fusion, Sam Altman is a major investor, and Microsoft already signed a 2023 agreement to buy Helion power as soon as 2028. The immediate takeaway isn’t “fusion is here”; it’s that top-tier AI demand is now big enough to seriously entertain exotic supply paths — and that alone changes how utilities, developers, and competitors think about long-dated capacity.

[Washington establishes transmission authority, tightens data center rules](#) in a two-part move: Senate Bill 6355 creates a **state electrical transmission authority**, and SB 5982 requires **independent generators feeding data centers** to comply with the state’s Clean Energy Transformation Act. Several more sweeping measures failed, but the direction of travel is still clear: states want more control over the grid build-out and less tolerance for “off-to-the-side” generation that doesn’t meet the same decarbonization obligations as everyone else. For developers, this isn’t just policy noise — it’s an early warning that interconnection strategy and behind-the-meter plans are going to be litigated in the legislature, not just engineered on the site plan.

In Europe, the compliance screws are tightening via the courts. [Environmental groups seek judicial review of CRU data centre decision](#) challenges Ireland’s December ruling that lets new data centres connect if, after six years, they supply as much electricity into the grid as they consume, with **80%** of that supply from renewables (and **20%** allowed from fossil fuels). The groups argue this breaches Ireland’s Climate Act and EU energy efficiency rules; the context is stark: data centres used **22% of Ireland’s electricity in 2024** and are forecast to reach **31% by 2034**. This is what the next phase of “license to operate” looks like in constrained markets — not just planning permission fights, but legal pressure on the grid connection framework itself.

Cooling hardware is also getting more aggressive — and more specific to AI racks. [Dell PowerCool eRDHx reduces AI data center cooling costs](#) introduces a closed-loop warm-water rear-door heat exchanger Dell says can cut cooling energy by **up to 74%** and support **up to 4x GPU density per rack**. Dell’s example is a hypothetical 1,024-server, 10MW deployment with **\$18.3M** in five-year energy savings, with availability listed as late 2025/April 2026. The point isn’t whether every site hits Dell’s modelled savings; it’s that OEMs are now selling cooling as a first-class performance lever for AI infrastructure, not a facilities afterthought.

States in the US Midwest are also moving to codify what data centers must disclose — and what the grid can build. [Nebraska proposes large-scale battery storage and data disclosures](#) would allow private electric suppliers to build **50-500MW** battery projects and even grant **eminent domain** for energy storage property. The same proposal includes provisions requiring data centers to file

annual reports on **electricity and water usage** and **tax exemptions**. Read together, it's a pragmatic package: build flexibility (storage) while forcing more daylight around the resource footprint — exactly the mix other states will copy as AI loads become too visible to ignore.

Behind the Headlines

Liquid cooling isn't just about hardware anymore — it's about operating models that make it scalable across a global footprint. [Ecolab and Salute partner to deliver CaaS for liquid cooling](#) plugs Ecolab's Cooling-as-a-Service into Salute's direct-to-chip liquid cooling operations service, bundling tools like 3D TRASAR monitoring, smart coolant distribution units, and “connected coolants” delivered via Ecolab's global service team. The subtext: as AI/HPC loops multiply, the constraint won't only be capex or facility design — it'll be day-2 operations, fluid quality, monitoring, and service coverage across dozens of sites. This partnership is a bet that cooling will be managed more like an outsourced utility layer, not a bespoke mechanical system every operator reinvents.

Security risk is also migrating to the control plane — and data center operators should treat that as a physical-world problem, not just a SOC ticket. [Attackers Target Management and Identity Planes in 2025 Network](#) flags a shift toward identity, supply chain, and management planes, with Cisco citing a **178% rise in device-compromise attacks** and nearly **40%** of top-targeted vulnerabilities hitting **end-of-life devices**. For operators, the uncomfortable implication is that “good perimeter plus patched servers” isn't the center of gravity anymore; out-of-band management, identity systems, and vendor dependencies are. If you run high-density AI halls with expensive, scarce hardware, the economic incentive for these attacks only increases.

Edge data centers keep showing up in places that don't look like traditional “edge” at all. [Aker BP and Armada deploy offshore modular data center](#) will put an Armada Galleon modular unit on an offshore rig on the Norwegian Continental Shelf to process drilling data in real time, positioning it as a reference installation for broader rollout and a standardised edge architecture (including local AI models). This is the industrial version of the AI build-out: the value of local inference isn't just latency — it's resilience and autonomy when

connectivity is constrained or operational stakes are high. Watch for more “reference installs” like this; they’re often the quiet start of a repeatable deployment playbook.

TELBO●RG