

Data Centre Briefing

May 02, 2026

Global

Key themes:

Google talks \$40bn Anthropic deal with 5GW compute; MARA buys 505MW Long Ridge gas plant for \$1.5bn; Seattle proposes 365-day moratorium amid 369MW data center requests; ERCOT/Texas SB 6 buffering pushes BESS microgrid controls

Google's reported plan to put up to \$40bn into Anthropic isn't just another AI funding headline — it's a capacity reservation dressed up as venture capital. In the same package, Google would commit as much as 5GW of TPU-based compute over five years, which is basically "multiple hyperscale campuses" worth of infrastructure. If that number holds, it's a loud signal that the AI arms race is now constrained less by model talent and more by land, interconnection queues, cooling, and steel-in-the-ground execution.

The Big Stories

[Google reportedly to invest up to \\$40B in Anthropic](#) would combine an initial \$10bn with up to \$30bn tied to performance — and, crucially, up to 5GW of compute supplied over five years. The structure effectively binds capital to physical delivery: Google isn't only betting on Anthropic's upside, it's locking in demand for its TPU stack and the data center buildouts required to make that promise real. The key watch item is whether "5GW" becomes the new unit of measure for frontier-model partnerships — and how quickly the grid and permitting reality pushes back.

That pushback is already visible inside the hyperscalers' own earnings numbers. [Hyperscalers boost data centre capex amid rising AI demand](#) flags strong cloud growth and heavy AI-related spend (AWS revenue \$37.6bn;

Google Cloud roughly \$20bn, up 63% YoY; Meta capex \$19.8bn), but also the bottleneck: power, cooling, and permitting. When Google is sitting on a backlog described as more than \$460bn and new capacity still takes 18–36 months to materialize, capital stops being the scarce input — deliverable megawatts do.

On the “if you can’t get power, buy power” end of the spectrum, [MARA to acquire Long Ridge Energy & Power for \\$1.5bn](#) is a notable bet on vertically integrated supply. MARA is buying a 505MW combined-cycle gas plant plus over 1,600 acres for a digital infrastructure campus, lifting its total capacity by about 65% to roughly 2.2GW, with an AI/IT buildout that targets 200MW online by mid-2028. It’s a reminder that merchant power + land is becoming a shortcut around interconnection queues — but also that “AI data center” strategies are increasingly being written in gas-turbine and combined-cycle terms.

Local politics is also hardening in ways that could turn siting into a first-order risk. [Seattle proposes 365-day moratorium on new data centers](#) would pause new large-scale builds while the city studies infrastructure, water, utility rates, land use, jobs, and public health; officials cite concerns that proposed projects could demand up to 369MW and strain Seattle City Light amid declining hydropower availability. This is not just a Seattle story — it fits the broader pattern documented in [Communities Push Back Against Rapid Data Center Expansion in US](#), from lawsuits and moratorium demands to outright political retaliation, with projects in the mix ranging from a proposed \$6bn, 360-acre site in Festus, Missouri to a Kevin O’Leary-backed Utah campus targeting up to 9GW. For investors, the message is simple: community consent is becoming as gating as interconnection.

Behind the Headlines

The industry’s biggest quiet problem right now is the widening gap between what gets announced and what can actually be delivered. [Hidden risks threaten deliverability of announced data center capacity](#) argues that failures to coordinate power, capital, land, engineering, and demand are piling up — and that “committed power, financing, execution, and customer pre-leasing” are what separate real projects from PowerPoint. This matters because the

market is starting to price “optional” capacity as if it’s inevitable, even while permitting timelines, transformer lead times, and construction sequencing say otherwise. If you’re underwriting anything greenfield, the hard question is no longer “is demand real?” — it’s “what exactly is the critical path, and who is contractually on the hook when it slips?”

Grid operators and regulators are also getting less tolerant of AI-era load behavior, which is pushing design toward buffering and controllability rather than raw megawatts. [BESS and microgrid control stabilize AI training data centers](#) lays out the case for integrating battery energy storage (BESS), microgrid supervisory control, and generation management software to smooth extreme, fast load swings — with ERCOT and Texas SB 6 cited as examples of emerging buffering requirements. The subtext: “behind-the-meter” is shifting from a resilience nice-to-have to a compliance pathway, especially where sites have leaned on expensive islanded gas-turbine generation. If this trend sticks, the winners won’t just be power developers — it’ll be controls, storage integration, and firms that can make a data center behave like a predictable grid participant.

Finally, there’s a growing split between regions that can credibly offer low-carbon, high-availability power at scale — and those that can’t — and the Nordics are leaning hard into that advantage. [Nordic blueprint for sustainable, scalable AI data center infrastructure](#) positions the region’s mix of near-zero carbon power, free-air and liquid cooling, and municipal heat reuse (including an Espoo data center supplying heat to Kesko) as a repeatable model for AI infrastructure. The Financial Times estimate cited — \$635bn of AI infrastructure spend this year across Google, Meta, Microsoft and Amazon — helps explain why “nice sustainability story” is turning into “competitive siting strategy.” The practical investor takeaway is that heat reuse, cooling choices, and grid carbon intensity are moving from ESG appendix to permitting and brand-risk fundamentals.