TELBORG

Data Center Briefing

March 20, 2026

Global

Key themes:
Google signs 1GW demand response with TVA, DTE and others; Project Cannoli: Google-DTE 1,600MW solar and 450MW storage; xAI Southaven Colossus 2 permitted 41 gas turbines up to 1.2GW; Fervo Cape Station closes $421m to deliver ~100MW Utah geothermal

Google just put demand flexibility on the same footing as generation: it's baked **1GW of demand response** into long-term contracts with multiple U.S. utilities, explicitly to speed up data centre interconnections by shifting ML workloads. That's a meaningful pivot from "find more megawatts" to "make the grid behave better." Pair it with Google's Michigan mega-solar-and-storage plan and you can see the playbook emerging: bring your own flexibility, show your work to regulators, and get to power faster than the next queue.

The Big Stories

Google integrates 1 GW demand response into long-term contracts with utilities including I&M, TVA, Entergy Arkansas, Minnesota Power, and DTE Energy, aiming to shift or reduce ML workloads when the grid is tight. The interesting bit isn't the headline capacity; it's the positioning: demand-side control as a grid-planning tool to connect data centres more rapidly, alongside Google's solar, geothermal, and long-duration storage efforts and initiatives like EPRI DCFlex. If this works, it gives hyperscalers a new lever in interconnection negotiations—and raises the bar for everyone else still showing up with a one-way load profile.

Google and DTE plan 1,600MW solar, 450MW storage data centre under "Project Cannoli," with a 20-year contract structure (CCAA) and a US$10m

community energy bill commitment. The numbers—**1,600MW solar** plus **450MW storage** (including **400MW/1,600MWh BESS** and **50MW LDES**)—signal just how far data-centre-led power procurement has moved from "offsetting" to effectively underwriting new grid assets. Put next to the 1GW demand-response move, Google is clearly trying to solve both sides of the bottleneck: build more supply *and* make its load schedulable.

xAI data center expansion raises pollution concerns in Mississippi-Tennessee after Mississippi regulators approved a permit for **41 permanent gas turbines** at the Southaven/Colossus 2 site—behind-the-meter generation estimated up to **1.2GW**. An independent study pegged potential health-related costs at up to **$44m per year**, with protests and legal challenges focused on emissions and permitting. This is the sharp edge of the "self-powering AI campus" trend: it can bypass some grid constraints, but it also drags data centres directly into local air-quality politics—and that's a slower, messier approval pathway than many developers seem to price in.

Fervo secures $421M non-recourse financing for Cape Station, funding Phase 1 construction of its enhanced geothermal project in Beaver County, Utah. The package includes construction-to-term debt plus bridges and letters of credit, targeting first power this year and roughly **100MW by early next year**, with an eventual path to **500MW** under PPAs (including Southern California Edison and Shell Energy). For data-centre investors, the takeaway is simple: bankable financing is starting to catch up to next-gen firm power, which matters because "always-on clean" is becoming a competitive differentiator in grid-constrained markets.

Tariff uncertainty and component shortages squeeze data center builds as revised U.S. tariffs face litigation and companies chase duty recoveries—while the bigger constraint is brutal: shortages of DRAM, HBM, SSDs and GPUs described as effectively **sold out through 2028**, with hyperscalers absorbing most inventory. This is a quiet reshaping of market structure: capital isn't the gating factor for many builds; supply chain priority is. If you don't have procurement leverage, your "AI-ready" timeline becomes a negotiating position, not a schedule.

Behind the Headlines

Liquid Cooling Transforms Data Center Operations, Not Just Technology is a useful reality check from T5: liquid cooling isn't a bolt-on; it forces changes across design, commissioning, and day-to-day operations. The operational details—fluid management, specialised training, and lifecycle continuity—are what will separate "we can host high density" from "we can run it reliably at scale." The most investor-relevant line is the implied timeline: T5 expects this to normalise over five years, which suggests a long runway of retrofit programs, new O&M skill stacks, and (inevitably) a few painful early failures.

Wixom drafts ordinance to limit data center impacts with proposed requirements around water use limits, cooling and decommissioning plans, noise/vibration studies, buffers, and—most importantly—financial guarantees to cover off-site infrastructure impacts. This is the "second wave" of U.S. municipal response: not trying to ban data centres outright, but forcing them to internalise costs that were previously socialised onto communities and utilities. The reference points—disputes in Michigan communities tied to an Oracle/OpenAI settlement in Saline Township and a paused Howell Township proposal—signal that local process risk is becoming repeatable and legible, which means it will be priced and litigated more systematically.

Suspicious Flags of Convenience Vessels Threaten Taiwan's Subsea Cables highlights rising incidents of suspicious maritime activity linked to undersea cable damage risks, with some vessels reportedly tied to China. For the data-centre industry, this isn't a niche telecom story; it's a reminder that "capacity" and "latency" are moot if physical connectivity becomes a geopolitical pressure point. Taiwan sits at the intersection of strategic compute, strategic semiconductors, and strategic communications—so cable resilience and monitoring are becoming part of infrastructure due diligence, not an afterthought.