

Data Center Briefing

March 18, 2026

Global

Key themes:

NVIDIA + AT&T/Comcast/Spectrum/T-Mobile AI Grid across 100,000 edge sites; Google Michigan data centre plan adds 2.7GW resources via DTE agreement; DOE directs FERC load interconnection rulemaking amid >2.2TW queue backlog; SK Group warns AI HBM wafer deficit >20% could persist until 2030

NVIDIA is trying to turn telcos into the next compute platform. At GTC, it and a roster of big operators pitched “distributed AI grids” that would convert roughly 100,000 network edge data centres into an inference fabric—and, in the most optimistic framing, scale to “more than 100GW” of AI capacity over time. If that sounds like a moonshot, it is—but it also lines up neatly with today’s other signal: power and interconnection are now the hard constraints, so everyone is shopping for new places to put compute (and new ways to run it).

The Big Stories

[Telecom Operators and NVIDIA Announce Distributed AI Grids](#) is the clearest attempt yet to reframe telco edge sites as an AI infrastructure layer, not just network plumbing. AT&T, Comcast, Spectrum and T-Mobile in the U.S., plus Indosat Ooredoo Hutchison in Indonesia, are lining up with NVIDIA (and partners like Cisco and HPE) to deploy RTX PRO 6000 Blackwell GPUs across distributed facilities for inference. The headline number—~100,000 distributed network data centres—matters because it’s the first credible “inventory” of sites that could absorb AI hardware without waiting years for new hyperscale campuses to clear power and permitting.

[Google to develop Michigan data center with energy-first plan](#) is a reminder that the hyperscalers are now designing projects around grid outcomes, not the other way around. Google is evaluating a site in Van Buren Township within DTE Energy's territory, and—crucially—its Clean Capacity Acceleration Agreement with DTE is framed around adding 2.7GW of new resources to the grid, plus a \$10m Energy Impact Fund, while Google covers electricity and infrastructure costs. The message to regulators and utilities is blunt: “We’ll fund the upgrades and bring supply—if you’ll move fast.”

[Interconnection queue hampers U.S. competitiveness; reforming grid interconnection urgently](#) puts numbers behind the bottleneck that's quietly setting the pace for U.S. digital infrastructure. The U.S. DOE has directed FERC to initiate rulemaking for load interconnection, while RMI argues generator interconnection reform must be accelerated and expanded—citing >2.2TW sitting in interconnection queues and an almost five-year average queue-to-operation timeline in 2024. The “why it matters” isn't just delay; it's cost volatility: RMI points to grid-enhancing technologies enabling 6.6GW in PJM and ~\$1bn/year savings at roughly \$0.1bn installation cost—exactly the kind of ROI that starts to look irresistible when every AI build is fighting for megawatts.

[US data centre pipeline hits 241 GW as growth slows](#) shows the U.S. market is still enormous—and getting more selective. Wood Mackenzie says disclosed pipeline reached 241GW by end-2025, but Q4 additions slowed to ~25GW, while ~183GW is tied to construction or supply agreements (about 22% of U.S. peak demand in 2025). Pair that with [US hyperscalers to spend \\$700B on AI-driven data centers](#): Moody's expects roughly \$700bn of capex this year from the six largest U.S. hyperscalers, rising to ~\$820bn by 2027, while warning about free cash flow strain, more debt, and builds being staged and contract-tied. Translation: the money is still pouring in, but boards are increasingly allergic to “build it and pray”—which will ripple through leasing, procurement, and who gets funded.

[SK Group warns memory wafer shortage could last to 2030](#) is the supply-chain counterweight to all the glossy capacity announcements. Chairman Chey Tae-won is warning of a >20% wafer deficit driven by AI HBM demand potentially persisting to 2030; the story also cites SK Hynix's 57% HBM and

32% DRAM market shares, and commentary that new fabs will be optimized for AI workloads, not the broader enterprise market. If you're underwriting AI data centre cashflows, this matters because "GPU availability" is only half the story—memory constraints (and pricing) can cap effective deployment and stretch commissioning schedules.

Behind the Headlines

[Data centre surge in Thailand raises water and air concerns](#) is what the next wave of market pushback looks like: not "data centres are ugly," but "data centres are competing for local resources." Thailand's expansion is described as more than 70 projects concentrated in the Eastern Economic Corridor (about 40 operational, ~20 under construction, ~10 planned), with the Board of Investment approving 36 projects worth ~US\$23bn in 2025 plus more valued at US\$3.1bn. Bridge Data Centres' 0.2GW QHI01 in Chonburi—backed by US\$2.8bn in bank financing and a 10-year water supply deal for ~3.3m m³/year—shows how quickly "utility-style" inputs like water contracts become central to bankability, not an afterthought.

[Microsoft tests MicroLED optical links for data center networking](#) is an underappreciated lever for AI economics: interconnect power. Microsoft researchers in Cambridge are testing MicroLED-based optical links using imaging fiber to push data across thousands of parallel channels, claiming ~50% lower interconnect power and lower costs, with commercialization targeted as early as 2027 (with partners including MediaTek). In a world where operators are already bumping into facility power envelopes, cutting network power isn't a nice-to-have—it's a way to convert "fixed megawatts" into more usable compute.

[CIG and ML&S launch \\$56M optical manufacturing joint venture](#) is a small deal with big implications: the optics supply chain is being rebuilt for AI fabrics. Cambridge Industries Group and Germany's ML&S are putting \$56m into a Dallas-headquartered JV to scale high-speed optical modules and Near-Packaged Optics in Mexico, explicitly targeting hyperscale data centre operators and AI infrastructure providers with 800G/1.6T interconnects. As racks densify and topologies get more bandwidth-hungry, the "where" of

optics manufacturing (and the ability to scale it) becomes a competitive input —right up there with power and land.

TELBO ● RG