

The Problem with Renewable-Only Power Procurement for AI Infrastructure

Brian McGoldrick | Founder, Sheephill Group

March 2026

At least one major hyperscaler is planning to power a 6+ GW AI campus in Texas using co-located solar, battery storage, and grid imports. ERCOT interconnection queue data reveals over 20 projects from a single renewable energy developer totaling roughly 11 GW of solar, wind, and battery capacity filed across West Texas and the Panhandle. The 1.8x overbuild ratio against stated datacenter demand is necessary because solar and wind cannot deliver firm power at nameplate capacity. But the deeper question is whether this architecture can deliver the 99.99% uptime that hyperscale AI infrastructure requires.

The Renewable Portfolio Model

The strategy is straightforward in concept: solar generates during peak daylight hours, battery storage dispatches during ERCOT's evening ramp (the duck curve), and the facility draws from the grid overnight when Panhandle wind drives prices negative or near-zero. The developer's entity structure in the ERCOT queue confirms they are building integrated renewable hubs at each site, pairing solar with wind and storage. The financial logic works on paper. Texas Panhandle wind capacity factors run 40–45%, overnight prices frequently go negative, and if blended correctly the all-in energy cost could be competitive while claiming carbon-neutral operations.

Where the Model Breaks Down: Reliability

Datacenter service level agreements require N+1 redundancy and 99.99%+ uptime. The failure modes of a renewable portfolio are correlated in ways that gas generation is not. A multi-day low-wind, cloudy period in the Panhandle, which occurs particularly in summer, would simultaneously reduce solar output, deplete battery reserves, and eliminate the cheap overnight wind imports the model depends on. At that point the facility is buying from the ERCOT grid at scarcity pricing, which has hit \$5,000/MWh during Texas summer peaks. That is not just a cost problem. It is an availability problem. ERCOT has experienced multiple grid stress events where reliable service to load was not guaranteed.

The Battery Supply Chain Problem

The battery cost decline narrative is almost entirely a story about Chinese manufacturing scale. CATL and BYD control roughly 55–60% of global lithium-ion cell production. Add other Chinese manufacturers and the figure exceeds 75%. Western-allied battery production from Samsung SDI, LG, SK, and Panasonic is subscale, more expensive, and largely committed to EV programs. For defense and intelligence-adjacent AI workloads, Chinese-sourced battery storage creates unacceptable supply chain risk. CFIUS review, ITAR compliance, and broader decoupling posture mean that battery storage from Chinese OEMs is either explicitly prohibited or operationally untenable at facilities processing sensitive workloads. The CHIPS Act logic extends to power infrastructure. Battery costs cannot continue to decline on the same trajectory if procurement is restricted to allied-nation suppliers.

Gas + CCS: The Operationally Superior Alternative

The honest comparison is not solar LCOE versus gas LCOE in isolation. It is the all-in cost of firm, dispatchable, 99.99% reliable power delivered to the datacenter fence. For renewables, that includes solar generation cost, battery cycling degradation, grid import costs including tail-risk scarcity pricing, the operational overhead of managing multiple generation sources, and whatever backup generation the facility quietly requires.

For behind-the-meter gas with carbon capture and sequestration, the equation is simpler: LCOE with 45Q tax credits, which can actually produce a net cost below standalone gas generation, plus a minimal grid connection for true backup. Gas turbines are manufactured by GE Vernova, Siemens Energy, and Mitsubishi Power, all allied-nation suppliers with deep existing supply chains. Delivery timelines are 2–3 years. Natural gas is not without delivery risk in Texas. Winter Storm Uri demonstrated that pipeline gas supply can be disrupted during extreme cold weather events. But that risk is acute and resolvable in days, not a permanent structural deficit that recurs daily. A behind-the-meter gas facility with on-site fuel storage and dual-fuel capability can manage weather-driven disruption. A renewable portfolio cannot engineer around the daily absence of sunlight. The technology simultaneously meets firm power requirements, N+1 redundancy, uptime SLAs, allied supply chain requirements, and carbon compliance.

What to Watch

The pattern is already established. Multiple hyperscalers building frontier AI infrastructure have chosen behind-the-meter gas generation first, with renewables layered on top. The projects that actually deliver firm power on a 2–3 year timeline to the highest-priority AI infrastructure will be gas. A renewable-heavy interconnection pipeline gives a hyperscaler

optionality and a carbon narrative, but when it comes to securing 99.99% uptime for a multi-gigawatt AI training campus, gas backup will likely be required regardless. Watch for hyperscalers or their renewable development partners to quietly file gas generation projects under new LLCs in the same counties where large renewable portfolios are already queued. That would confirm the renewable-only model is not sufficient.

Analysis based on ERCOT GIS Report (December 2025) interconnection queue data cross-referenced with public hyperscaler datacenter project announcements.