To: Prospective Clients
From: Derek Tan Date: Aug 14, 2026
Subject: The Global Energy Renaissance

Semiconductors (“chips”) are the new hottest trade in AI. At the time of writing, the VanEck Semiconductors ETF (NASDAQ:SMH) is up ~60% YTD, far outperforming all other sectors in the S&P 500 by a wide margin. For reference, the S&P 500 is only up ~8% YTD. A combination of accelerating AI-driven capex, supply constraints, earnings revisions, and strategic demand from hyperscalers/governments resulted in a market rally that has become so feverish, the traditional Mag 7 has been effectively replaced by the Parabolic 7 (SanDisk, Marvell, Micron, Intel, Dell, AMD, Broadcom), named after the shape of their recent share price growth.

However, beyond the first-order effects of this chip-driven rally, a more fundamental physical constraint on AI's growth is emerging. This analysis will examine one critical question: whether chips are the critical binding constraint to AI’s growth, or whether the physical power systems and infrastructure are the true bottleneck to unlocking the next stage of innovation.

From first principles, semiconductors remain to be the obvious “picks-and-shovels” bottleneck. Hyperscalers (Microsoft, Alphabet, Amazon, Meta) have pledged to spend ~$700B in AI infrastructure capex in 2026, a significant amount of which would be directed towards buying thousands of new chips. Using NVIDIA’s earnings as a proxy, with FY’26 revenues up 65% YoY with gross margins of ~75%, given that NVIDIA is already the largest player and still growing at such a rapid pace, it suggests that the broader AI hardware ecosystem is probably still undersupplied.

Yet the physical constraints on AI's growth run deeper than silicon. According to the IEA, global data center electricity consumption is set to nearly double from ~485 TWh in 2025 to ~950 TWh by 2030, with AI-specific facility consumption tripling over the same period after already surging ~50% in 2025 alone. Hyperscalers can order chips on a matter of quarters; they cannot order gigawatts of new power on the same timeline. U.S. grid interconnection queues now hold over 2,300 GW of proposed generation and storage — more than the country’s entire installed capacity — with wait times stretching 5–7 years, and in some constrained markets, a decade or more. Compounding this, persistent shortages in transformers, gas turbines, and switchgear are pushing out construction timelines and costs for new generation, regardless of how quickly chip supply normalizes.

This asymmetry — chip lead times measured in quarters versus power infrastructure lead times measured in years — is why we believe power, energy, and physical infrastructure, not semiconductors, are the true binding constraint on AI’s next leg of growth.

Hyperscalers’ pivot toward direct power procurement (e.g., Microsoft’s $15.2B UAE commitment, Meta’s $10B Louisiana campus, and dedicated wind PPAs) is a tacit admission of this reality. From an investor’s perspective, this reframes where durable value in the AI trade accrues from the chip layer toward utilities, grid equipment makers, independent power producers, and on-site generation providers positioned to solve the physical bottleneck.

Sources: Company filings