AI, Power, and Strategic Infrastructure

The U.S. Grid Buildout, Global Competition, and an Investor Framework

Prepared for research and discussion | September 2026

Abstract

Artificial intelligence is converting dependable, deliverable electricity into a strategic input for economic output, frontier computing, industrial competitiveness, and national security. This paper examines the emerging power-and-infrastructure investment cycle through the linked lenses of U.S. grid constraints, data-center demand, the U.S.-China strategic competition, household affordability, and public-equity selection. It argues that the central scarcity is not electricity in the abstract but power delivered at the right location, voltage, reliability standard, and date. Consequently, the most durable opportunities may lie across a system of bottlenecks—grid equipment, engineering and construction capability, regulated network assets, firm generation, fuel delivery, and grid-enhancing technologies—rather than in a single headline AI trade. The paper concludes with a set of research questions and a practical underwriting program centered on nodes, bottlenecks, contracts, and downside scenarios.

Keywords: artificial intelligence; data centers; electricity demand; transmission; utilities; energy infrastructure; national competitiveness; U.S.–China strategic competition; public equities.

1. Introduction

Artificial intelligence is increasing the strategic value of reliable electric power. Modern AI systems depend on data centers that require unusually large, concentrated, and reliability-sensitive loads, thereby linking technological expansion to the availability of generation, transmission, substations, electrical equipment, cooling, and fuel supply. The resulting investment opportunity is not adequately described by the phrase “buy electricity.” Rather, AI has intensified a broader infrastructure problem: the ability to deliver power at the right grid node, at the required voltage and uptime, and on the timetable demanded by major compute projects. The U.S. Department of Energy’s 2026 National Transmission Needs Study identifies accelerating load growth from data centers, AI, and other large loads as a major source of transmission-system need. This development changes the analytical task for investors. The relevant question is not merely whether power demand rises, but who owns, builds, supplies, finances, regulates, and ultimately pays for the assets required to serve that demand.

2. The Power-and-Delivery Constraint

The power thesis begins with the distinction between generating energy and delivering dependable capacity. A new data-center campus can require hundreds of megawatts or more than a gigawatt of service, while its operator may need high uptime, redundant supply, rapid energization, and a credible expansion path. Those requirements stress local substations, high-voltage transmission, distribution feeders, generation reserves, natural-gas delivery systems, backup capacity, and interconnection processes simultaneously. The International Energy Agency projects continued rapid global electricity-demand growth through 2030 and identifies data centers among the principal drivers of the U.S. increase. Yet a national demand forecast can obscure the localized nature of the constraint. Power may be abundant in one region but unavailable at the relevant node, and a project can be economically attractive while still waiting years for a transformer, line upgrade, generation connection, or permitting decision.

3. Why the Grid Is the Core Asset

Transmission and distribution infrastructure are central because they translate generation into usable economic capacity. DOE estimates that U.S. transmission congestion imposed approximately $11 billion in costs during 2023, with a large share arising during a relatively small number of high-stress hours. This highlights a critical feature of the AI-power cycle: average annual energy is not enough when the system must remain reliable under extreme weather, equipment outages, and coincident peak demand. Investments in transmission corridors, substations, transformers, switchgear, conductors, protection systems, advanced power-flow controls, dynamic line ratings, reconductoring, and voltage uprates may therefore create value in more than one macroeconomic outcome. They can support AI load, electrification, reliability, renewable integration, industrial expansion, and replacement of aging assets. From an investor’s standpoint, assets that address multiple demand scenarios generally deserve a higher-quality designation than assets whose returns depend on a single AI spending forecast.

4. AI as Strategic Capability

AI should be understood as an enabling capability for economic productivity, cyber defense, intelligence, logistics, scientific research, industrial design, and military decision support. Its power requirements do not make electricity the only determinant of national capability, because semiconductor access, communications, software, data, talent, capital, and institutional capacity remain essential. Nevertheless, reliable power increasingly functions as a prerequisite for scaling advanced compute. PwC’s Global Data Centre Outlook estimates a central-case total of $31.6 trillion in global data-center capital expenditure through 2050, with the outcome varying materially according to AI adoption, chip availability, and geopolitical conditions. The outlook also emphasizes that data-center investment is recurring rather than strictly one-time because servers, networking, storage, and other ICT equipment must be periodically refreshed. This transforms the issue from a discrete construction boom into a potentially multi-decade requirement for power-system capacity, supply-chain resilience, and deployment speed.

5. America, China, and Execution

The U.S.–China competition should not be reduced to a contest over installed gigawatts. China retains major advantages in centralized coordination, manufacturing depth, clean-energy deployment, storage investment, and the ability to construct infrastructure at scale. The United States retains powerful advantages in advanced semiconductors, hyperscale cloud platforms, venture and public capital, research institutions, software ecosystems, and many frontier AI developers. The IEA expects China to account for close to half of global electricity-demand growth through 2030, illustrating the scale of its industrial and power-system challenge. The U.S. opportunity is not to imitate China’s political structure, but to shorten the interval between credible demand and reliable energized capacity. That requires faster interconnection, clearer cost allocation, more predictable permitting, greater equipment availability, skilled labor, and better regional planning. For investors, time-to-powered-operation may become a more useful strategic metric than announced capital expenditure.

6. Consequences for Households and Politics

For households, the buildout offers both benefits and risks. New data centers and grid projects can create construction, engineering, skilled-trade, manufacturing, and operations employment; broaden tax bases; and potentially improve reliability if investment is allocated efficiently. However, the same projects can increase electricity bills if costly network and generation upgrades are socialized while large customers receive favorable tariffs or if demand forecasts later fail to materialize. The IEA has emphasized affordability as a central electricity-policy issue, noting that household electricity prices in many countries have increased faster than incomes since 2019. The decisive policy question is thus one of cost allocation: whether hyperscalers and developers make sufficient contributions to interconnection and network upgrades, whether they accept minimum-bill or take-or-pay commitments, and whether they provide flexible-load capability when the system is stressed. The political legitimacy of the AI buildout will influence rate cases, local permitting, data-center moratoria, and the realized returns of regulated utilities.

7. Investment Framework

A robust portfolio approach should diversify across the physical bottlenecks instead of relying on a single narrative. Grid-equipment and electrification providers can benefit from demand for transformers, switchgear, protection, power quality, and substation modernization. Engineering and construction firms can monetize the constrained capability required to execute transmission, distribution, power-generation, and data-center mechanical and electrical projects. Regulated utilities can earn on rate-base expansion when commissions permit prudent capital investment and allocate costs fairly, while independent power producers may benefit where firm capacity and energy are scarce. Gas and fuel infrastructure can provide balancing and reliability value, particularly where new loads need dispatchable supply before long-lead transmission or nuclear assets are available. Early-stage technologies such as small modular nuclear, distributed fuel-cell generation, and certain connectivity plays may offer asymmetric upside, but should be sized as higher-risk options rather than treated as substitutes for businesses with proven earnings and durable backlogs.

Table 1. Illustrative Public-Market Exposure Map

Layer	Illustrative companies	Potential mechanism	Core underwriting risk
Grid equipment and electrification	ETN, HUBB, POWL, GEV	Substations, distribution, protection, power quality, grid modernization	Cycle normalization; order delays; new capacity; valuation
Engineering and construction	PWR, FIX, MTZ	Transmission, utility, mechanical and electrical project execution	Labor, materials, margin pressure, backlog quality
Regulated utilities	NEE and regional utilities	Rate-base expansion and large-load service	Rate cases, affordability, financing, cost allocation
Firm generation	CEG, VST	Reliable generation and capacity value in constrained markets	Power-price volatility, policy, outages, hedging
Fuel infrastructure	LNG, ET, NFG and peers	Fuel delivery and dispatchable-power support	Commodity exposure, permitting, contracts, policy
Higher-risk optionality	OKLO, SMR, BE, LITE, LUMN	On-site generation, advanced nuclear, optical and connectivity pathways	Technology, funding, dilution, commercialization

Note: Illustrative tickers are not investment recommendations. Company-specific valuation, balance-sheet strength, backlog conversion, customer concentration, regulatory exposure, and geographic asset positioning require independent review.

8. Research Questions for Investors

8.1 What portion of announced data-center load is real?

Investors should calculate probability-weighted deliverable load rather than treating all announced gigawatts as equivalent. A project should be tracked from land control and utility service agreement through transmission-upgrade responsibility, equipment procurement, financing, construction start, and energization. This framework distinguishes proposals from bankable projects and is likely more informative than national headline demand estimates. The key output is a regional project ledger that identifies credible megawatts by balancing authority, customer quality, and expected in-service date.

8.2 How flexible is AI demand?

The load shape of AI may prove more consequential than total annual electricity consumption. Training workloads can potentially move across time or geography more readily than latency-sensitive inference, while government, health, financial, and industrial workloads may demand tighter reliability and location constraints. If operators can curtail or shift consumption, non-firm interconnection, storage, advanced grid controls, and demand-response contracts could stretch existing infrastructure further. If they cannot, the investment case for firm generation, on-site power, and expensive network reinforcement strengthens.

8.3 Who bears the cost?

The answer differs by state, tariff, commission, utility structure, and market design. Investors should examine contribution-in-aid-of-construction requirements, minimum bills, take-or-pay obligations, demand charges, upgrade-cost responsibility, stranded-cost protections, and customer exit rights. A utility may face extraordinary load growth yet create limited shareholder value if regulators socialize upside while leaving the firm with financing, construction, and demand-risk exposure. Conversely, carefully structured large-load tariffs can protect households and support durable utility investment.

8.4 What is the binding bottleneck?

The relevant constraint may be high-voltage transmission, transformer supply, substation capacity, turbine lead times, skilled labor, gas transport, rights-of-way, or local permitting. Each bottleneck has a different time-to-relief, competitive structure, and set of public-market beneficiaries. Investors should map equipment lead times, manufacturing capacity, domestic sourcing, backlog duration, labor availability, and pass-through mechanisms. This process converts a broad infrastructure theme into testable, company-level underwriting assumptions.

8.5 How does geopolitics change the geographic map?

A more fragmented technology environment could reduce some cross-border compute demand while increasing domestic and sovereign workloads. PwC's scenario analysis suggests that chip restrictions can lower overall data-center capex, whereas a data-sovereignty world may redistribute capacity across regions rather than eliminate it. Investors should distinguish export-sensitive training workloads from locally required inference, regulated data, defense, and critical-industry deployments. The latter may be more resilient to geopolitical fragmentation but could require localized power, communications, and secure infrastructure.

8.6 Can U.S. institutions build quickly without waste?

The central policy challenge is to accelerate essential investment while avoiding expensive overbuild if AI monetization, hardware efficiency, or project completion diverge from optimistic forecasts. A resilient plan uses phased interconnections, flexible-load contracts, standardized equipment, grid-enhancing technologies, credible demand milestones, and scenario-based transmission planning. Firms with assets useful across high, medium, and lower-demand outcomes should be more robust than firms dependent on a single extreme forecast. The investor should treat execution speed, regulatory clarity, and contracting discipline as economic moats rather than as secondary details.

9. Practical Research Program

A practical program starts with geography, not tickers. Select several relevant power markets—such as ERCOT, PJM, MISO, Southeastern utility territories, and selected Western systems—and build a localized map of load forecasts, planned generation, retirements, reserve margins, transmission approvals, congestion, gas constraints, and large-load tariffs. Then link each prospective investment to a specific node and bottleneck: a contractor's awarded transmission backlog, an equipment supplier's transformer or switchgear capacity, a utility's approved rate-base plan, or a generator's contract coverage and fleet reliability. Test each holding under at least three cases: accelerated AI demand, delayed interconnection with moderate demand, and a weaker AI monetization outcome. Monitor load-forecast revisions, queue status, customer contracts, equipment lead times, backlog conversion, cancellation rates, tariff dockets, capacity markets, reserve margins, and balance-sheet financing. The goal is not to predict the most dramatic AI future, but to identify companies whose economics remain attractive across a range of plausible grid and demand outcomes.

10. Conclusion

AI has elevated electricity infrastructure from a conventional utility topic to a central question of economic competitiveness, technological sovereignty, national security, and household affordability. The durable investment opportunity lies in the gap between announced compute demand and the physical, regulatory, and contractual machinery required to energize it reliably. The U.S. can retain and extend its advantages in frontier AI without replicating China’s system, but it must materially improve the speed and coordination of its grid-development process. Investors should resist the temptation to treat all AI-adjacent infrastructure equities as interchangeable beneficiaries. The highest-quality analysis focuses on deliverable power, regional constraints, contractual cost allocation, backlog quality, regulatory treatment, and scenario resilience. In this framework, the decisive question is not whether AI needs power; it unquestionably does. The decisive question is which firms control the scarce assets and capabilities that transform that need into reliable, profitable infrastructure.

Selected Public Sources

International Energy Agency (2026). *Electricity 2026: Executive Summary*. Paris: IEA.

International Energy Agency (2026). *Key Questions on Energy and AI*. Paris: IEA.

U.S. Department of Energy (2026). *National Transmission Needs Study: Draft for Consultation and Public Comment*. Washington, DC: DOE Office of Electricity.

PwC (2026). *Global Data Centre Outlook 2026–50: Where $31.6 trillion of capex flows in the era-defining AI build-out*.

Brookings Institution (2026). *How will the United States and China power the AI race?*.

Ember (2026). *China Energy Transition Review 2025*.