CARF WORKING PAPER

A Comparative Analysis of Global AI Energy Infrastructure

Demand, Constraints, and Policy Responses

Author: Ozgun Tutuncuoglu

THE NUMBERS AT A GLANCE

448 TWh	945 TWh	$5.2 trillion	2030–2032
World data center electricity (2025) — equal to Italy's entire consumption	Projected by 2030 — equal to all of Japan's electricity today	McKinsey base-case infrastructure investment needed by 2030	Earliest realistic date for first commercial small nuclear reactors

SIX KEY TAKEAWAYS

1. AI is now a major electricity consumer — and growing faster than any other sector. The world's data centers used ~448 TWh in 2025. That number is on track to more than double by 2030, driven not just by humans using chatbots, but by AI systems talking autonomously to each other — what researchers call 'agentic AI.'
2. The United States leads in AI models but its electricity grid cannot keep pace. New data centers are built in 12–24 months; connecting them to the grid now takes up to 10 years. The most underappreciated constraint is not capital or technology — it is a shortage of licensed electricians.
3. China has a structural energy advantage that may matter as much as its chip disadvantage. Central planning allows energy infrastructure to be built before it is needed. China added more electricity generation capacity in 2024 alone than one-third of the entire US power grid.
4. Europe is falling further behind. Grid connection approvals in major data center hubs take up to 13 years. Regulatory ring-fencing is expected to raise consumer electricity prices 10–20% in affected regions. Without urgent reform, the EU cannot close the gap.
5. Small modular nuclear reactors (SMRs) are the only scalable, continuous, low-carbon answer for AI's power needs — but not before 2030–2032. Solar and wind are intermittent; AI systems need uninterrupted power.
6. The energy infrastructure bottleneck is not a distant risk — it is already materialising. As of April 2026, 30–50% of US data centers planned for 2026 face delays or cancellations. If this gap widens while AI-related equities remain priced for unimpeded expansion, the result could be a cascading financial crisis: AI agent demand fuelling inflation, a sharp repricing of technology assets (beginning with Nasdaq-listed AI names), spreading into financial sector balance sheets and eventually triggering a recession. Markets are currently ignoring this risk.

Author: Ozgun Tutuncuoglu | Full academic working paper follows | April 2026 | CARF Review Draft

ABSTRACT

The rapid expansion of artificial intelligence (AI) infrastructure is generating unprecedented demand for electrical power, constituting one of the most significant structural shifts in global energy systems since large-scale industrialization. This paper provides a comparative institutional analysis of how the United States, the People's Republic of China, and the European Union are responding to this challenge across five dimensions: energy demand trajectories, grid infrastructure capacity, nuclear and renewable energy policy, regulatory frameworks, and supply-side constraints.

Drawing on 2025 institutional data from the IEA, Gartner, Goldman Sachs Research, McKinsey Global Institute, Deloitte, and BloombergNEF, the paper establishes that global data center electricity consumption reached approximately 448 TWh in 2025 and will rise to 945 TWh by 2030 under the IEA base case. The paper's original contribution is a multi-source forecast reconciliation and sensitivity analysis framework that resolves divergent institutional projections — including BNEF's 8.6% US share estimate, IEA's 9–12% range, and Deloitte's 123 GW capacity forecast — and identifies the five parameters most responsible for forecast heterogeneity, with agentic AI adoption rate carrying the largest uncertainty range (±150–230 TWh in the US alone by 2030).

The analysis identifies five principal findings: (1) agentic machine-to-machine AI traffic is a structurally undermodeled demand accelerant; (2) China's centrally directed model provides a meaningful structural advantage over market-led US and EU approaches; (3) the EU faces the most acute infrastructure constraints and risks continued relative decline, exacerbated by ring-fencing policies that will raise consumer electricity prices 10–20%; (4) Small Modular Reactors are the only viable pathway to scalable, continuous, low-carbon baseload power for AI but will not be commercially available before 2030–2032; and (5) a shortage of skilled electrical workers — not capital or technology — is the binding near-term constraint in all three jurisdictions. The paper concludes with an author's original risk assessment: the energy infrastructure bottleneck, if left unaddressed, may trigger a cascading financial crisis — beginning with AI asset repricing and spreading through inflation, recession, and broader market instability — the severity of which markets are currently failing to price.

Keywords: *Artificial Intelligence; Energy Infrastructure; Data Centers; Small Modular Reactors; Grid Capacity; Agentic AI; Forecast Sensitivity; Labor Constraints; Ring-Fencing; Systemic Risk; US-China-EU Comparison*

TABLE OF CONTENTS

M.2 Agentic AI Demand: Assumption Transparency

The paper's treatment of agentic AI — autonomous machine-to-machine interaction as a demand multiplier — rests on three explicit assumptions that peer reviewers should evaluate:

- Current IEA and Gartner base-case models project AI energy demand based primarily on historical and near-term human query rates scaled by forecast AI penetration rates. Agentic workflows — in which an AI model autonomously executes multi-step tasks involving web search, code execution, API calls, and coordination with other AI models — consume computational resources estimated by Goldman Sachs Research (2025) at one to three orders of magnitude larger than single human queries, depending on task complexity.
- No institutional model as of April 2026 has published a validated, data-grounded estimate of the ratio of agentic to human-initiated AI traffic for 2028–2035. The Goldman Sachs (2025) assertion that agentic traffic 'may dwarf human-generated traffic by the late 2020s' is a qualitative directional claim, not a quantified projection. This paper treats it as upside risk supporting the IEA Lift-Off scenario (1,300 TWh by 2030) rather than adjusting the central estimate.

METHODOLOGICAL NOTE	*This is identified as the primary area where additional empirical grounding would most strengthen the analysis. Future versions should incorporate operational data from deployed agentic AI systems when hyperscalers publish consumption disclosures.*

1. Global AI Energy Demand: Scale, Scenarios, and Emerging Drivers

1.1 Current Scale and Baseline Projections

Gartner's November 2025 report estimates global data center electricity consumption at approximately 448 TWh in 2025, an 8% increase over the 415 TWh recorded in 2024. The IEA's April 2025 special report 'Energy and AI' projects this rising to 945 TWh by 2030 and approximately 1,200 TWh by 2035 under its base case. AI-optimized server electricity consumption is projected to grow nearly fivefold between 2025 and 2030, from 93 TWh to 432 TWh — rising from 21% to 44% of total data center electricity (Gartner, 2025). Goldman Sachs Research (2025) estimates the 2023-to-2030 growth trajectory at 220%, equivalent to adding a new top-ten power-consuming nation to the global grid within a decade.

Figure 1.1a: Global Data Center Electricity Consumption Projection, 2025–2035

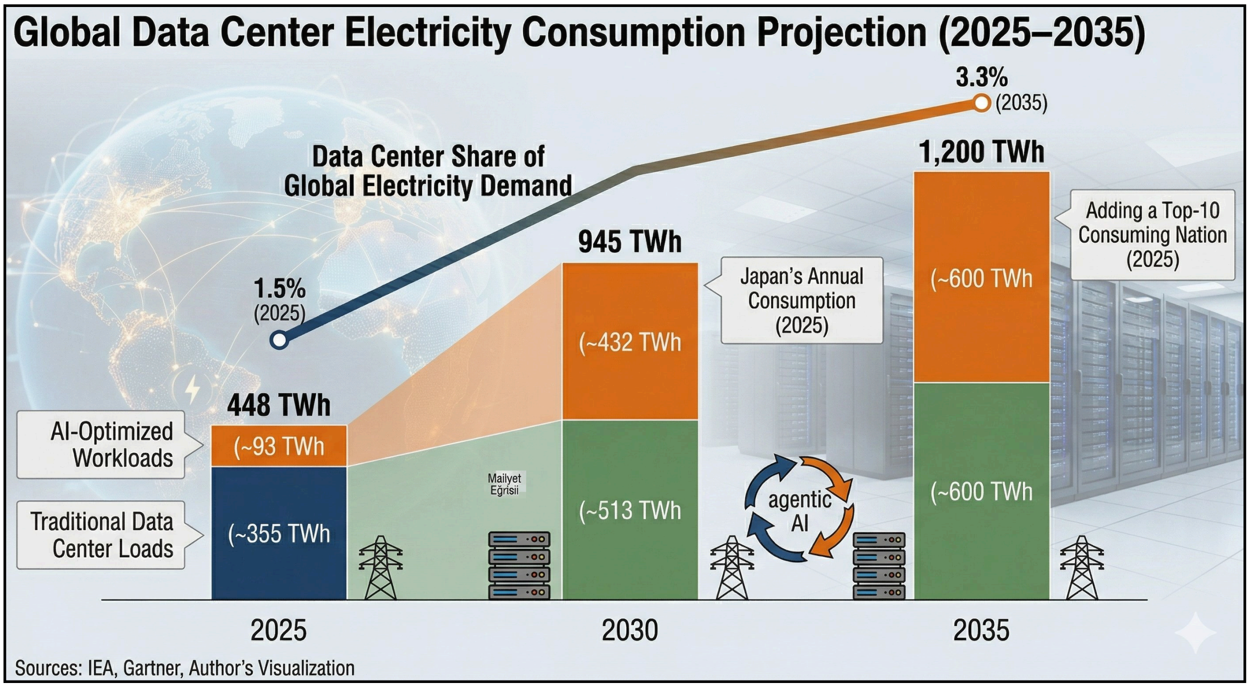

Sources: IEA, Gartner, Author's Visualization

Figure 1.1a. Sources: IEA (2025); Gartner (2025); Author's visualization.

Figure 1.1b: Global Data Center Electricity Demand — Growth Trajectory and Scenario Range

Year	Headwinds (Low)	Base Case	High Efficiency	Lift-Off (High)	Context / Comparator
2023 (actual)	~415 TWh	~415 TWh	~415 TWh	~415 TWh	Baseline year (IEA 2024)
2025	~430 TWh	~448 TWh	~440 TWh	~465 TWh	Gartner estimate; ~Italy's total electricity
2027 (proj.)	~520 TWh	~620 TWh	~570 TWh	~720 TWh	First SMR demos expected; agentic AI scaling
2030 (proj.)	~650 TWh	~945 TWh	~800 TWh	~1,300 TWh	Base case ≈ Japan's entire electricity
2035 (proj.)	~700 TWh	~1,200 TWh	~970 TWh	~1,700 TWh	High case > Germany + France combined

Figure 1.1. Source: IEA Energy and AI (April 2025). All projections except 2023/2025 are scenario-dependent estimates. Lift-Off scenario assumes agentic AI at scale by 2028–2030.

1.2 Regional Concentration of Demand

The US and China together account for approximately 80% of global data center electricity consumption growth through 2030. Figure 1.2 shows the estimated 2025 regional distribution. Note: China's official figures likely understate actual consumption.

Figure 1.2: Regional Data Center Electricity Consumption, 2025

Region	2025 Consumption (TWh)	Share of Global Total	Key Driver
United States	~210 TWh	46.9%	5,426 data centers; dominant hyperscaler presence
China	~140 TWh	31.3%	Centralized planning; eastern coastal concentration; coal-dominated
EU + United Kingdom	~78 TWh	17.4%	Fragmented; market share declining from 25% (2015) to 15% (2024)
Japan	~17 TWh	3.8%	Mature market; efficiency-focused
South Korea + Taiwan	~8 TWh	1.8%	Semiconductor-driven; export-oriented
Southeast Asia	~5 TWh	1.1%	Fast-growing; Singapore/Malaysia hubs
Rest of World	~5 TWh	1.1%	Emerging markets; minimal current share
TOTAL	~448 TWh	100%	≈ 1.5% of global electricity consumption

Figure 1.2. Regional data center electricity consumption, 2025 estimates. US and China together account for ~78% of global consumption. Sources: IEA Energy and AI (2025); Gartner (2025). China figures may be understated due to limited official disclosure.

1.3 Forecast Reconciliation and Sensitivity Analysis

Published forecasts of US data center electricity demand vary substantially across institutions. This section makes an original analytical contribution: systematic reconciliation of four major forecasts and a sensitivity framework identifying which parameter assumptions drive most of the divergence.

1.3.1 Cross-Institutional Forecast Comparison

Institution	Metric	US 2030 Forecast	US 2035 Forecast	Key Definition Difference
IEA (April 2025)	TWh/year	~426 TWh (~9–12% of US demand)	Not specified	Covers all data centers; 2024 baseline 183 TWh; percentage applied to IEA's 2030 US demand projection (~4,670 TWh)
BloombergNEF (July 2025)	% of US electricity	~8.6% (≈410 TWh implied)	8.6% of larger 2035 grid	BNEF anchors to 2035 as reference year; applies to a higher projected total US demand (~4,770 TWh). TWh figure is consistent with IEA.

Institution	Metric	US 2030 Forecast	US 2035 Forecast	Key Definition Difference
Deloitte (Dec. 2025)	Installed GW capacity	~78 GW (2030)	~123 GW (2035)	Capacity (GW) ≠ consumption (TWh). 123 GW × 8,760 hrs × ~50% utilization ≈ 539 TWh — consistent with IEA/ BNEF when utilization rates are reconciled.
Goldman Sachs (2025)	% of US electricity	~8% (≈380 TWh)	Not specified	Conservative utilization; excludes some edge/enterprise facilities. Directionally consistent with IEA/BNEF base case.

Table 1.2. Cross-Institutional US Data Center Demand Forecast Comparison. Author's reconciliation. Absolute TWh projections from IEA, BNEF, and Goldman Sachs are within a 12% range — a high level of agreement for ten-year forecasts.

1.3.2 Sensitivity Analysis Framework

Parameter	Central Assumption	Low Case Effect	High Case Effect	2030 US Impact Range
AI hardware efficiency improvement rate	15% annual reduction in energy per computation	−10% on total DC energy (if 25% p.a.)	No change (if 5% p.a.)	±90 TWh
Agentic AI adoption rate (machine-to-machine)	Nascent in 2030; <5% of AI compute	−30 TWh (if delayed to 2033+)	Up to +200 TWh (if dominant by 2029)	±150–230 TWh — highest-impact variable
Grid interconnection speed (regulatory reform)	PJM queue persists; partial FERC reform	−15 TWh (worsens)	Up to +60 TWh (fast-track adopted)	±75 TWh
US total electricity demand trajectory	~4,670 TWh total US by 2030	Lower denominator → higher % share	Higher denominator → lower % share	±0.8 percentage points on % share metric
Capacity utilization (Deloitte reconciliation)	~48% average utilization	45% → 485 TWh for 123 GW fleet	55% → 592 TWh for 123 GW fleet	±53 TWh for Deloitte 123 GW figure

Table 1.3. Sensitivity Analysis. Author's original analysis based on IEA (2025), BNEF (2025), Deloitte (2025), Goldman Sachs (2025). Agentic AI adoption rate is the dominant uncertainty parameter.

RESEARCH FINDING	*Agentic AI is the single parameter with the largest potential impact on 2030 forecast outcomes (±150–230 TWh in the US alone), yet no institutional model has published a validated agentic-to-human traffic ratio. This is the primary unresolved uncertainty in the projection landscape and the key variable determining whether the IEA base case or Lift-Off scenario prevails.*

Figure 1.2a: Global Data Center Energy Source Transformation — 2025 vs. 2035

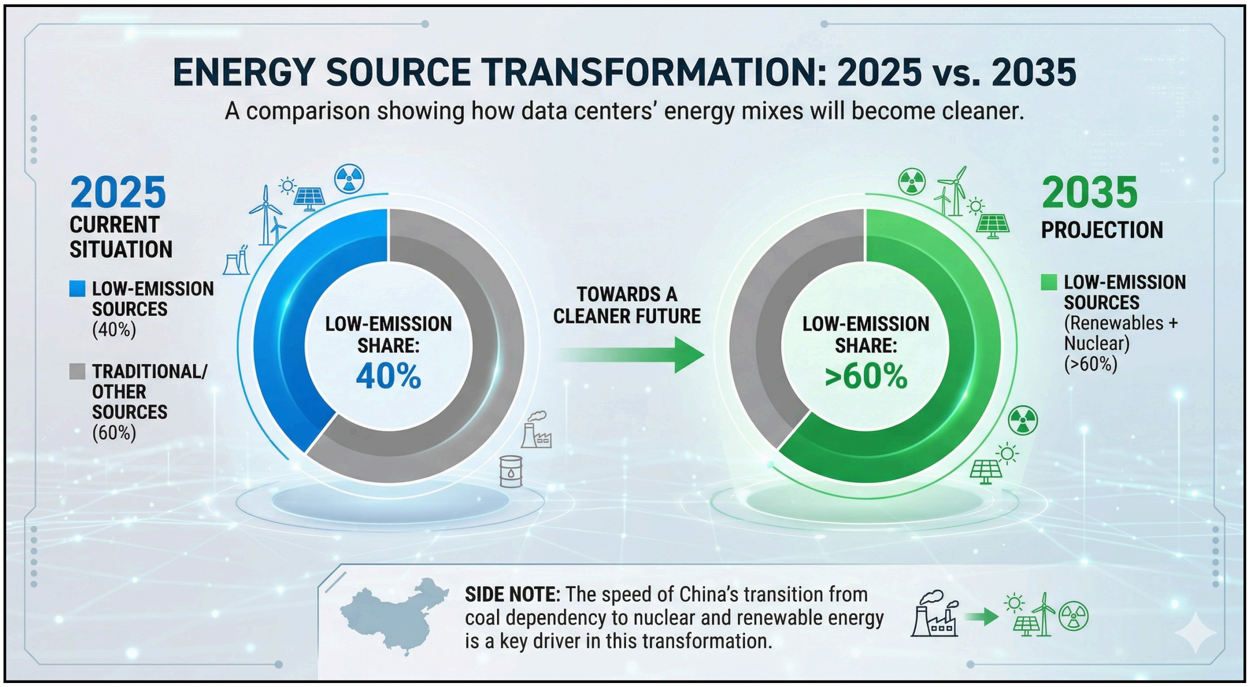

Figure 1.2a. Low-emission sources projected to rise from 40% to over 60% of data center electricity supply by 2035. Source: IEA (2025); Author's visualization.

2. Regional Infrastructure Comparison: United States, China, and the EU

2.1 United States: Scale Leadership and Grid Bottlenecks

The United States hosts 5,426 operational data centers (March 2025) and consumed approximately 210 TWh for data center operations in 2025 — roughly 5% of national electricity. The IEA projects this rising to 9–12% by 2030, with data centers accounting for nearly half of all US electricity demand growth through that period. The central constraint is not generation capacity but grid delivery: PJM Interconnection — the regional transmission organization serving 13 states including Pennsylvania, New Jersey, and Maryland, covering approximately 65 million consumers — faces a waiting list of new connection requests extending beyond 2035. Gas turbine order books at GE Vernova, Siemens Energy, and Mitsubishi are full through 2029. High-voltage transformer lead times have extended to 2.5–4 years globally.

RESEARCH FINDING	*Goldman Sachs Research (June 2025) identifies licensed electricians and high-voltage technicians — not capital or technology — as the single most immediate near-term constraint on US AI energy infrastructure expansion. The US requires 200,000+ additional workers in these trades. Apprenticeship programs require 4–5 years to produce qualified journeymen; the existing pipeline is structurally insufficient to meet projected demand through 2030.*

2.2 China: Centrally Directed Energy Expansion

China added more new electricity generation capacity in 2024 than one-third of the entire installed US grid. Its centrally directed planning model allows energy infrastructure to be constructed in anticipation of demand — producing a structural speed advantage over market-led approaches. Chinese data centers consumed approximately 140 TWh in 2025, with coal accounting for ~70% of supply due to the geographic concentration of facilities in eastern coastal provinces. By 2035, IEA projects renewables and nuclear will supply ~60% of China's data center electricity. The 'Eastern Data Western Compute' (EWCRT) initiative uses Ultra High Voltage transmission to migrate compute loads to renewable-rich western provinces. A geopolitical asymmetry defines competition: the US controls access to advanced AI semiconductors; China controls the clean energy manufacturing supply chains (solar panels, batteries, rare earth processing) essential to data center power infrastructure.

2.3 European Union: Regulatory Ambition and Infrastructure Deficit

The EU's global data center market share fell from 25% (2015) to 15% (2024). IEA projects EU data center capacity growing only 70% by 2030 under current policy — far below the Commission's goal of tripling capacity. Grid connection approvals reach 13 years in major hubs; Dublin and Amsterdam have imposed moratoriums.

Country/Region	2030 DC Demand (TWh)	Structural Advantage	Primary Constraint
France	~30 TWh	Nuclear-dominated grid (~70%); competitive electricity prices; €112B AI investment announced 2025	Administrative permitting timelines
Germany	~35 TWh	Engineering industrial base; central location	North-south grid congestion; EnEfG compliance costs

Country/Region	2030 DC Demand (TWh)	Structural Advantage	Primary Constraint
United Kingdom (UK)	~20 TWh	Regulatory flexibility; active Great British Nuclear SMR program	High construction costs; limited carbon-free baseload
Nordic Countries	~22 TWh	EU's lowest-cost clean electricity; hydro baseload; fastest permitting	Distance from major data center demand centers
Ireland / Netherlands	Capped	Established hyperscale hub; transatlantic subsea connectivity	Grid saturation; active connection moratoriums
Poland / Spain	Emerging fast growth	Lower land and energy costs; available grid capacity	Less mature data center ecosystem

Table 2.1. EU/UK Data Center Market Analysis. Sources: IEA (2025); European Commission (2025); Eurelectric (2025).

Figure 2.1: Regional Data Center Infrastructure Comparison — Bottlenecks and Strategic Advantages

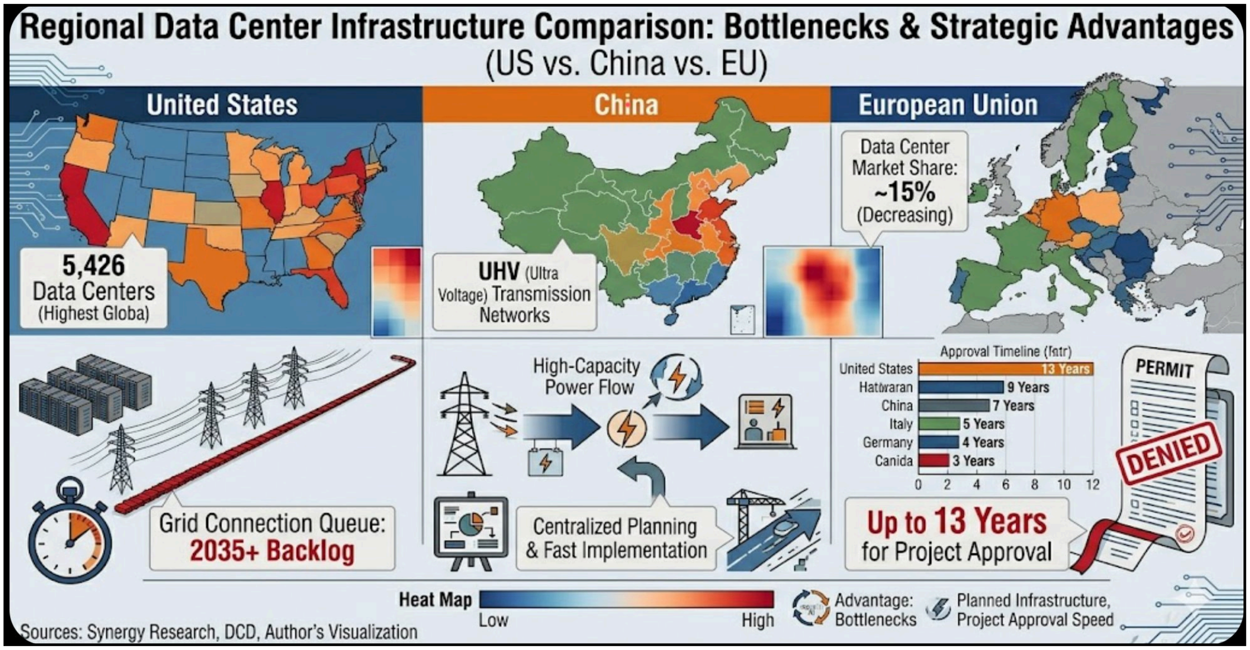

Figure 2.1. US grid connection queue extends to 2035+; China's centralized UHV networks enable fast deployment; EU grid approval timelines reach up to 13 years. Sources: Synergy Research; DCD; Author's visualization.

RISK WARNING	*Ring-fencing regulations — requiring AI operators to fund the full costs of their own grid infrastructure — are expected to raise consumer electricity prices by 10–20% in high-concentration markets including parts of Virginia (US), Ireland, and Germany. Separately, geopolitical risk from Chinese dominance of rare earth element (REE) processing — essential for wind turbine magnets, EV batteries, and transformer components — creates a supply chain dependency that could constrain the renewable buildout required to power AI data centers in the US and EU. Any export restriction on Chinese REEs (analogous to 2023 gallium/germanium controls) would materially slow clean energy deployment timelines.*

3. Small Modular Reactors: Technology and Environmental Assessment

3.1 Technology Definition and AI Infrastructure Rationale

A Small Modular Reactor (SMR) is a nuclear fission unit designed for factory fabrication and modular site assembly, typically generating 50–300 MW(e) per module. The rationale for AI infrastructure is straightforward: AI training and inference operate continuously at near-maximum power draw for weeks or months; solar PV output falls 60–80% under cloud cover; battery storage at 1–5 GW campus scale for multi-day periods is not economically viable; natural gas creates carbon liabilities incompatible with net-zero commitments. SMRs are the only currently available pathway to scalable, continuous, low-carbon baseload power at the densities required by frontier AI infrastructure. Note: nuclear energy generates radioactive waste and should be described as 'low-carbon' or 'carbon-free in operation' — not 'clean energy,' a term that incorrectly implies the absence of any environmental impact.

3.2 SMR Power Module: Structure and Safety Architecture

Figure 3.1: NuScale SMR Power Module — Cross-Section Schematic

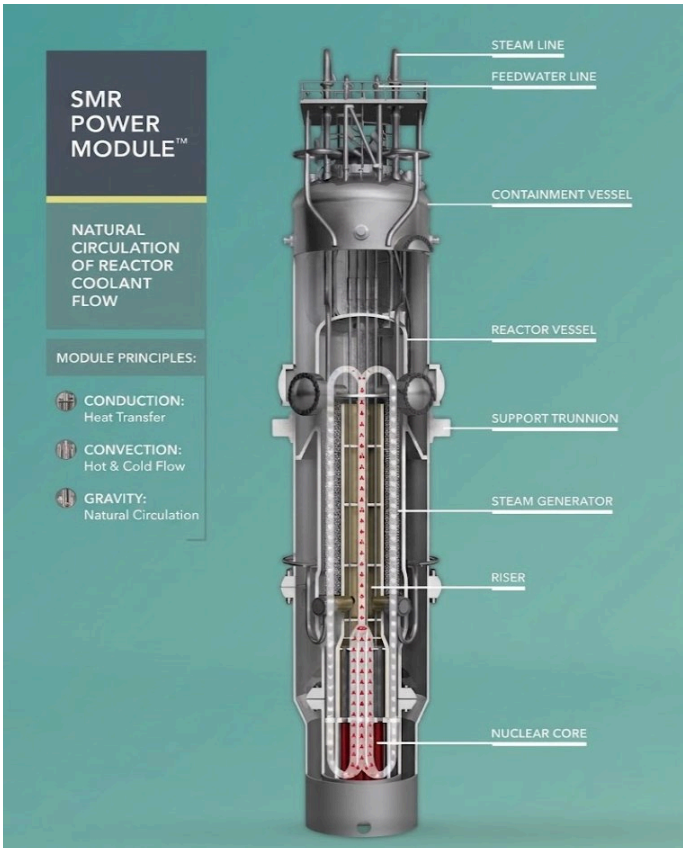

Figure 3.1. SMR Power Module cross-section. Natural circulation of reactor coolant via conduction, convection and gravity — no active pumps, external power or operator intervention required. The nuclear core heats coolant that rises through the riser, transfers heat to the steam generator, and circulates back by natural convection inside the containment vessel. Source: NuScale Power (2025).

3.3 Construction Challenges and Commercialization Timeline

Challenge	Explanation	Current Mitigation Path
High first-of-kind cost	$12,000–$15,000/kW installed (~3× a gas plant); reflects design finalization, bespoke components, regulatory unfamiliarity	Serial production target: $4,000–$6,000/kW. Long-term technology company PPAs provide demand-side signal.
Regulatory timelines	NRC design certification historically 5–10 years; EU frameworks vary by member state	Trump executive orders (May 2025) directed NRC to accelerate SMR licensing. NuScale Standard Design Approval issued May 2025.
No OECD commercial precedent	No SMR has yet operated commercially in any OECD country; ±50% cost uncertainty and ±3–5 year timeline uncertainty	China's ACP-100 'Linglong One' (Hainan Province) expected to connect to grid in 2026 — first commercial benchmark.
Fuel supply chain	Some advanced designs require HALEU (High-Assay Low-Enriched Uranium), not yet at commercial scale in Western supply chains	Amazon's $500M X-energy investment (2024) includes HALEU fabrication. US DoE enrichment programs active.
Nuclear waste management	Radioactive waste requires isolation for thousands of years; permanent geological disposal frameworks incomplete	SMR designs using existing spent fuel as input reduce volumes. US geological repository legislation progressing.

Table 3.1. Sources: IEA (2025); NRC (2025); IAEA (2025); DoE (2025).

3.4 Environmental Profile

Dimension	SMR Performance	Comparative Benchmark
Lifecycle carbon intensity	4–12 g CO_2-equivalent per kWh (full lifecycle)	Gas: 490 g/kWh; Solar PV: 20–50 g/kWh; Wind: 7–15 g/kWh; Coal: 820 g/kWh. Nuclear among lowest-carbon sources (IPCC AR6, 2022).
Radioactive waste volume	Substantially lower per unit energy than large conventional reactors; some designs use existing spent fuel	All US nuclear waste since 1958 ≈ 10,000 m^3 — smaller than a standard residential subdivision
Water consumption	Design-dependent; advanced gas-cooled and molten salt designs use significantly less than pressurized water reactors	Conventional nuclear: ~2.5 L/kWh; open-cycle gas: ~0.5 L/kWh
Land use	500 MW SMR campus: typically 2–5 hectares operational area	500 MW solar: ~700–1,000 hectares; 500 MW wind: 50–100 km^2 with turbine spacing
Accident safety	Passive safety systems — no pumps, external power, or operator action needed for emergency core cooling	Extends Gen III+ passive safety (AP1000); underground siting reduces surface-level accident risk

Table 3.2. Sources: IPCC AR6 (2022); IAEA (2025); NEA/OECD (2023); IEA (2025).

Period	Key Milestones	Jurisdictions
2025–2027	NRC licensing acceleration; TerraPower Natrium (WY) construction; Kairos Hermes 2 demo (TN); China ACP-100 Linglong One grid connection (Hainan, 2026); UK Great British Nuclear competition awards	US, China, UK
2027–2030	First demo reactors achieve criticality; Google/Kairos commercial power; TerraPower first output; tech operator PPAs begin delivering	US, China
2030–2032	First-of-kind commercial SMR operational in US; project financing enabled; cost reduction curve initiated	US, Canada, China
2032–2035	Serial production achieves cost targets; SMR becomes standard for large AI campus power; EU first commercial SMR (France or UK)	US, EU, Canada, China

Table 3.3. Sources: IEA (2025); NRC (2025); TerraPower (2025); Kairos Power (2025); IAEA (2025).

Figure 3.2: SMR Commercialization Roadmap — Meeting AI's Demand for Reliable 24/7 Baseload Power

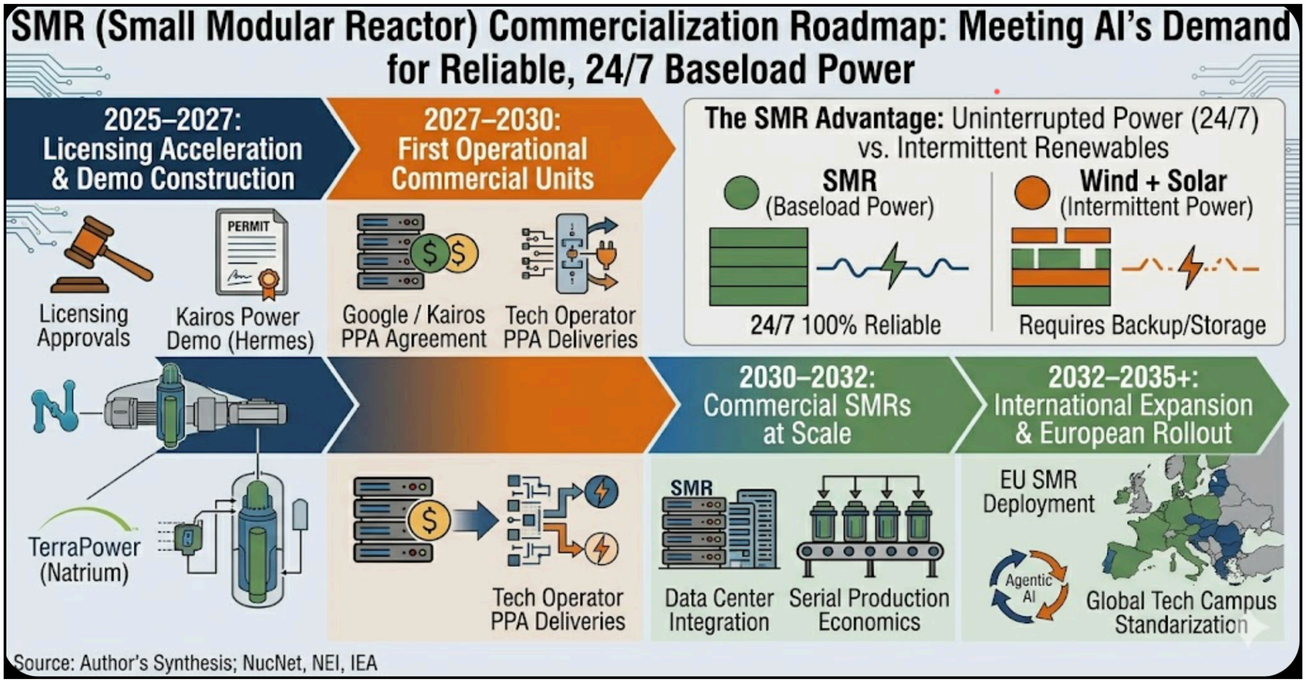

Figure 3.2. SMR deployment phases: 2025-2027 licensing and demo construction; 2027-2030 first operational commercial units; 2030-2032 commercial scale; 2032-2035 international expansion. Sources: NucNet; NEI; IEA; Author's Synthesis.

4. Infrastructure Constraints, Labor Bottlenecks, and the 2026 Deployment Crisis

---

4.1 The Infrastructure Lead-Time Mismatch

Infrastructure Type	Typical Build Time	US Status (2026)	EU Status (2026)
Data center construction	12–24 months	No significant delay	12–24 months
Grid connection (new substation)	3–10+ years	PJM queue to 2035; FERC reform partial	3–13 years; Dublin/Amsterdam moratoriums
Natural gas CCGT plant	3–5 years	Turbine order books full to 2029	4–6 years
High-voltage transformer	2.5–4 year lead time	Global constraint; eases ~2027	Same global constraint
Transmission line (overhead)	4–8 years	Local opposition; regulatory fragmentation	4–10 years; cross-border complexity
Existing nuclear plant restart	1–3 years	2–3 plants feasible (Crane/TMI, Palisades)	1–2 plants (France only near-term)
SMR — first-of-kind	8–12 years from decision	2030–2032 earliest	2033+ earliest
SMR — serial production	5–7 years post-first-of-kind	Post-2032	Post-2035

Table 4.1. Sources: IEA (2025); FERC (2025); S&P Global (2025); NRC (2025).

RESEARCH FINDING	*Goldman Sachs Research's 'Six P's' framework identifies 'People' — skilled electrical workers — as the most immediate constraint, ahead of equipment or capital. The US needs 200,000+ additional licensed electricians and grid technicians; the 4–5 year apprenticeship pipeline makes this structurally impossible to resolve before 2028–2029. This bottleneck is not addressable through capital expenditure or policy streamlining alone.*

Figure 4.1: Goldman Sachs Research — 'The Six P's' Bottleneck Analysis for AI Infrastructure

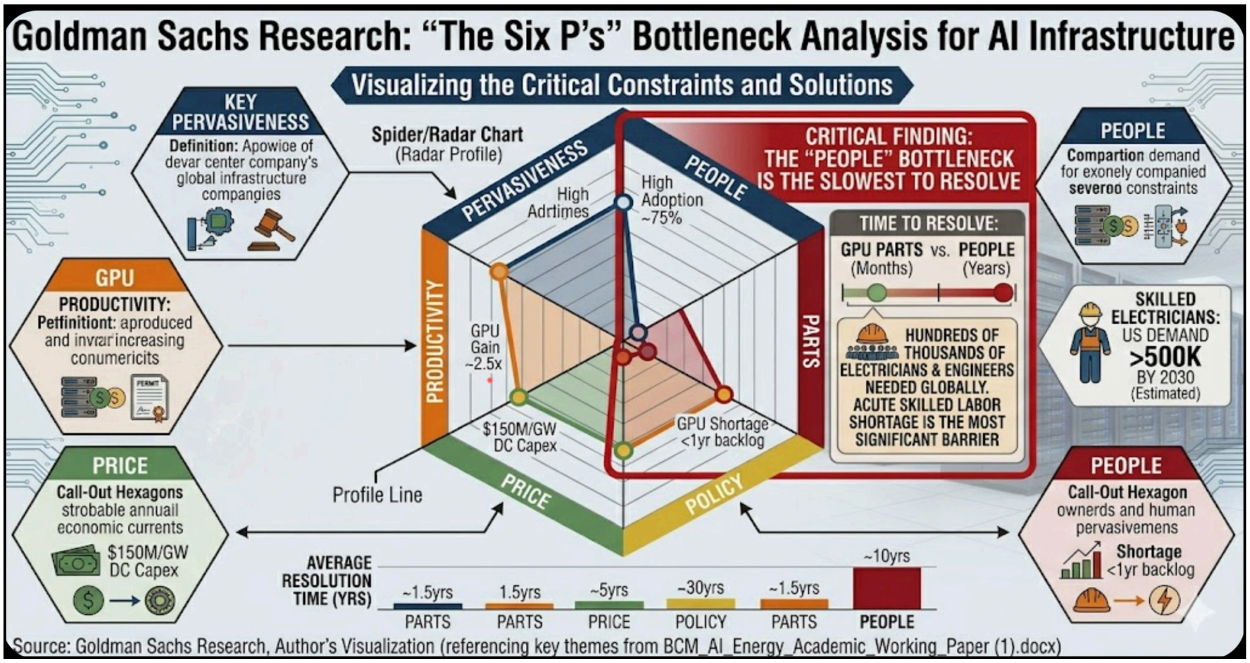

Source: Goldman Sachs Research, Author's Visualization (referencing key themes from BCM_AI_Energy_Academic_Working_Paper (1).docx)

Figure 4.1. The 'People' bottleneck (skilled electricians/engineers) requires up to 10 years to resolve — far longer than any equipment or price constraint. GPU parts: <1 year; Policy: ~30 years (generational); People: ~10 years. Source: Goldman Sachs Research (2025); Author's visualization.

4.2 Capital Requirements

Investment Category	2025–2030 Estimate (Base)	Primary Sources
Data center construction and real estate	~$1.2 trillion in new asset value	JLL Global DC Outlook 2026
IT equipment (processors, networking, storage)	$1.0T – $2.0T	JLL (2026); McKinsey (2025)
Energy infrastructure (generation + transmission)	~$1.3T (~25% of McKinsey total)	McKinsey (2025); S&P Global (2025)
US utility grid capex (47 investor-owned utilities)	>$1.0T cumulative 2025–2029	S&P Global / Regulatory Research Associates (2025)
Nuclear investments (SMR + existing plants)	~$150B through 2030	IEA (2025); DoE (2025)
TOTAL (McKinsey base case)	$5.2T ($3.7T – $7.9T range)	McKinsey Global Institute (May 2025)

Table 4.2. Sources as indicated. MAG7 companies (Alphabet, Amazon, Apple, Meta, Microsoft, NVIDIA, Tesla) collectively spent approximately $500–600B on AI infrastructure in 2024–2025; this is projected to exceed $1 trillion annually by 2028–2030 (Goldman Sachs, 2025).

4.3 Live Evidence: The 2026 US Data Center Delay Crisis

The constraints documented in Sections 4.1 and 4.2 are no longer theoretical. As of April 2026, the energy and supply chain bottlenecks have produced a measurable and accelerating deployment crisis in the world's largest AI infrastructure market.

Metric	Finding	Source
2026 announced US DC capacity	~16 GW across 140 construction projects slated for 2026 delivery	Sightline Climate / Bloomberg, April 2026
Actually under construction	~5 GW (~31% of announced pipeline)	Sightline Climate, April 2026
Proportion delayed or cancelled	30–50% of 2026 pipeline	Sightline Climate / Bloomberg, April 2026
Expected 2026 online capacity	~12 GW (vs. 16 GW announced)	Sightline Climate, April 2026
2027 announced pipeline	~21.5 GW announced	Sightline Climate, April 2026
2027 under active construction	~6.3 GW (~29% of announced)	Sightline Climate, April 2026
2025 delay precedent	26% of 2025 capacity slipped; 10% pushed back COD without notice	Sightline Climate, 2025
Chinese transformer imports	Surged from ~1,500 units (2022) to ~8,000 units (2025 through Oct.)	Wood Mackenzie / Bloomberg, 2026
Chinese share of US battery imports	>40% of US battery imports; ~30% of some transformer/switchgear categories	Bloomberg, April 2026

Table 4.3. US Data Center Deployment Reality Check, April 2026. Sources: Sightline Climate Data Center Outlook (February 2026 and April 2026); Bloomberg (April 2026); Wood Mackenzie (2026).

RISK WARNING

The gap between announced pipeline (16 GW) and active construction (5 GW) for 2026 US data centers is directly attributable to the constraints documented in this paper: transformer shortages, switchgear delivery backlogs, skilled labor deficits, and Chinese supply chain dependency. Critically, capital is not the constraint — MAG7 companies are collectively spending over $650B in 2026 alone. The bottleneck is entirely physical and regulatory. This real-time evidence validates the paper's core thesis and, by extension, strengthens the financial risk assessment in Section 5.5: if 30–50% of the physical infrastructure underpinning AI equity valuations is delayed by 1–3 years, the earnings assumptions currently embedded in AI sector valuations are materially at risk.

5. Conclusions: Energy Bottleneck, Geopolitical Risks, and the Financial Crisis Risk

5.1 Principal Findings

- AI energy demand is structurally large and likely to exceed current base-case projections, primarily due to the agentic AI paradigm — the highest-uncertainty variable in current institutional forecasts, with a ±150–230 TWh impact on the 2030 US estimate alone.
- Cross-institutional forecast divergences (IEA, BNEF, Deloitte, Goldman Sachs) are largely reconcilable through metric-definition adjustments; absolute TWh projections are within a 12% range, representing high agreement for ten-year infrastructure forecasts.
- China's centrally directed energy model provides a structural speed advantage over market-led US and EU approaches, compounded by its dominance of clean energy manufacturing supply chains.
- Skilled labor scarcity is the most immediate and underappreciated constraint. Unlike equipment or capital, the electrician pipeline cannot be rapidly scaled through spending increases.
- SMRs are the only viable pathway to scalable, continuous, low-carbon baseload power at AI data center densities — but commercial availability before 2030 is not realistic. First-of-kind: 2030–2032; serial production scale: 2032–2035.
- Ring-fencing regulations are expected to raise consumer electricity prices 10–20% in affected markets. Chinese REE dominance creates a geopolitical supply chain dependency that could materially slow clean energy deployment in the US and EU.

5.2 AI Model Adoption Growth: A Demand Foundation

To contextualize the energy demand trajectory, it is important to establish the underlying driver: the explosive and accelerating adoption of AI models by users and enterprises globally. The following chart illustrates the trajectory of AI assistant adoption and the projected forward path to 2030 and 2035.

Figure 5.1: Global AI Assistant Monthly Active Users — Historical and Projected

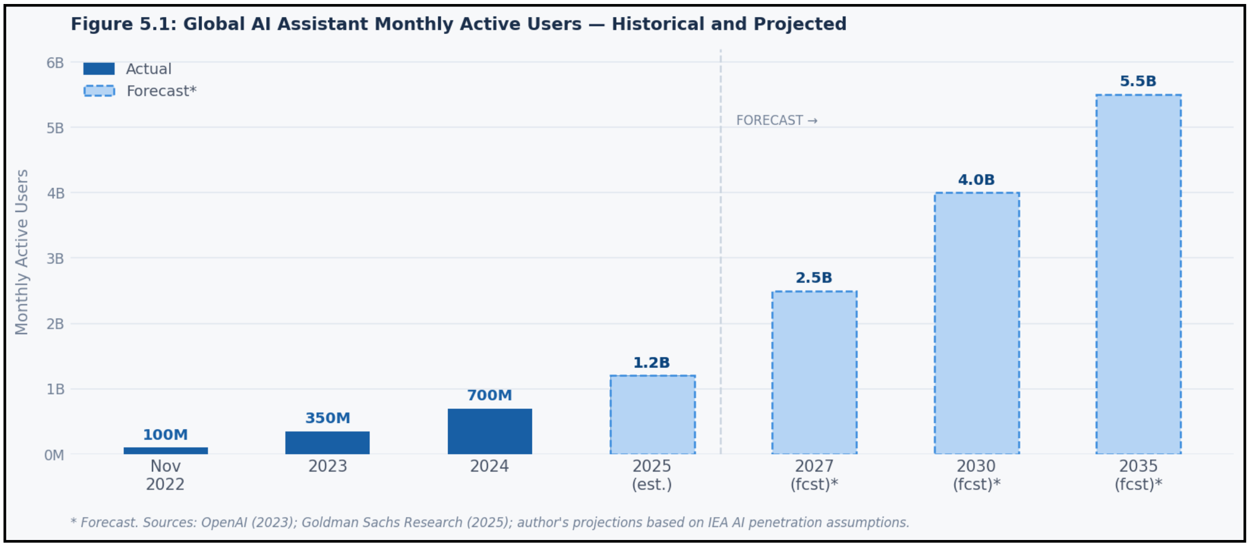

*Figure 5.1. * Forecast. Actual data through 2025 (est.); forecast 2027–2035 based on IEA AI penetration assumptions and Goldman Sachs Research (2025). Human-initiated AI user growth drives baseline energy demand; agentic AI traffic (Figure 5.2) is projected to be the dominant compute driver by the late 2020s.*

5.3 Agentic AI: The Multiplier Effect

Human user growth, while impressive, may be the smaller component of future AI energy demand. The true demand multiplier is agentic AI — autonomous AI systems executing multi-step tasks without human intervention in each step. A single agentic workflow (web search + document analysis + code generation + API calls + validation) can consume 100–1,000× the computational resources of a single human chat message. The following chart shows the projected trajectory of deployed AI agents:

Figure 5.2: Deployed Agentic AI Systems — Historical and Projected (Global Instances)

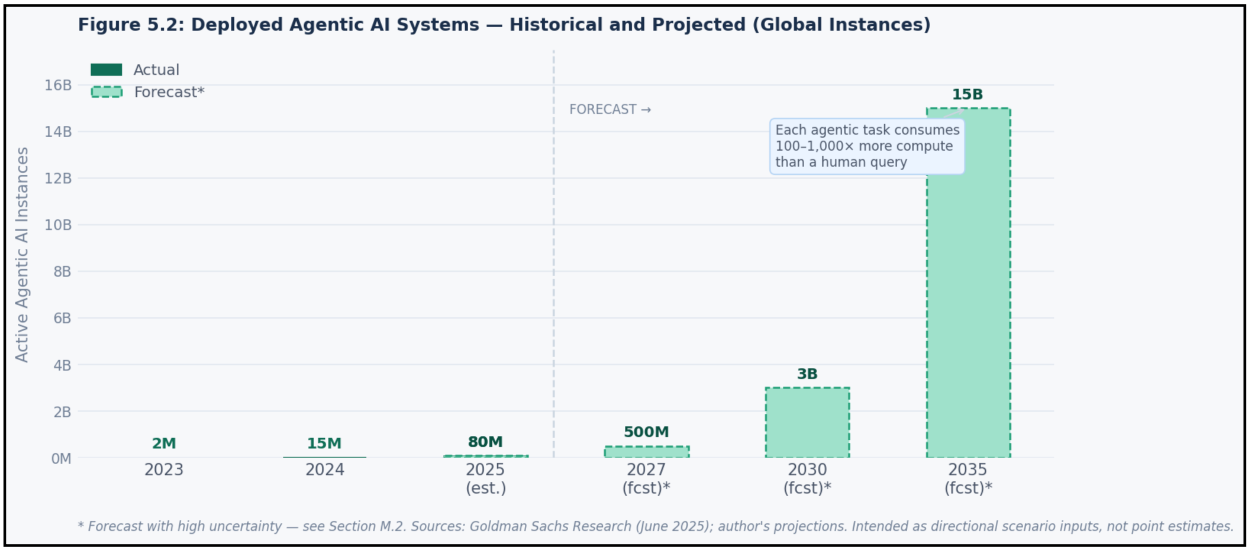

*Figure 5.2. * Forecast with high uncertainty — see Section M.2 for assumption transparency. Each agentic task consumes an estimated 100–1,000× the compute of a single human query. If agentic AI scales as projected, current IEA Lift-Off scenarios (1,300 TWh by 2030) may prove conservative. Sources: Goldman Sachs Research (June 2025); author's projections.*

5.4 The MAG7 Spending Trajectory and the Infrastructure Gap

The Magnificent Seven technology companies — Alphabet, Amazon, Apple, Meta, Microsoft, NVIDIA, and Tesla — collectively invested approximately $500–600 billion in AI infrastructure, research, and development in 2024–2025. Goldman Sachs Research projects this figure exceeding $1 trillion annually by 2028–2030 as enterprise AI adoption accelerates and agentic workloads scale. Microsoft alone committed $80 billion in FY2025; Amazon exceeded $75 billion; Google and Meta together allocated approximately $175 billion.

The central problem is structural: technology capital deploys in months; energy infrastructure takes years to a decade. A hyperscaler can commission and fill a data center in 18–24 months; the grid connection required to power it at full load may not be available for 8–10 years in constrained markets. This creates a compounding mismatch that worsens as AI adoption accelerates. The bottleneck is not a shortage of capital, and it is not — at the fundamental level — a shortage of technology. The barriers are regulatory (interconnection queue timelines, permitting), physical (transformer lead times, transmission construction), and, most acutely, human (the licensed electrician shortage).

RISK WARNING

The energy infrastructure gap is widening faster than it is being closed. Even if all announced policy reforms (FERC interconnection streamlining, SMR licensing acceleration, EU fast-track permitting) are implemented immediately and successfully, the 2028–2030 period will feature acute power constraints for data center deployment in large parts of the US and EU. This is not a hypothetical risk — it is already visible in hyperscaler behind-the-meter gas turbine deployments, utility interconnection moratoriums, and data center siting migration to power-available regions.

5.5 Systemic Financial Risk: The Energy Bottleneck as Crisis Catalyst

The energy infrastructure gap poses a systemic financial risk that the mainstream AI investment literature has not adequately analysed. The transmission mechanism is as follows: the AI sector has attracted capital on the premise of rapid, unimpeded infrastructure deployment. As of April 2026, however, Sightline Climate and Bloomberg report that 30–50% of US data centers planned for 2026 have been delayed or cancelled — not due to a lack of capital or technology, but due to transformer shortages, switchgear supply constraints, and skilled labor deficits that no amount of spending can resolve in the near term. This is not a hypothetical risk. It is already occurring.

The financial crisis transmission chain operates through four sequential channels. First, as AI deployment delays become apparent, hyperscaler earnings guidance faces downward revision — capital expenditure of $600B+ annually generates returns only when data centers are activated and monetized. Second, Nasdaq-listed AI infrastructure names (NVIDIA, Microsoft, Alphabet, Amazon, Meta) — which collectively represent over 25% of S&P 500 market capitalisation — face repricing as revenue timelines extend. Third, this equity correction spreads into financial sector balance sheets: banks with leveraged AI infrastructure loans, data center REITs, private equity funds with AI capex exposure, and sovereign wealth funds with concentrated technology positions all face mark-to-market losses. Fourth, the broader macroeconomic feedback loop: AI agent demand inflation (compute, energy, skilled labor all simultaneously scarce) compounds general price pressures, eroding real incomes and consumer demand in ways that interact with existing debt cycles to produce recessionary dynamics.

RISK WARNING

RISK WARNING — AUTHOR'S ORIGINAL ASSESSMENT: The energy infrastructure bottleneck represents, in this author's judgment, an unprecedented and currently underpriced systemic financial risk. The mechanism is distinct from any prior financial crisis: it is not leverage (2008), not overvaluation of intangible assets (2000), and not a currency crisis (1997). It is a physical constraint on the monetisation of real assets — data centers and GPU clusters — whose valuation is already embedded in global equity indices. As AI agent demand accelerates and the deployment gap widens, the probability of a disorderly repricing event rises. Bond markets, currently pricing AI infrastructure bonds at investment-grade spreads, have not adequately modelled the scenario in which deployment timelines slip by 2–3 years. Equity markets, pricing MAG7 names at 25–35× forward earnings, are implicitly assuming unimpeded AI monetisation. Both assumptions are increasingly difficult to sustain given the physical realities documented in this paper. This author's assessment: the probability of a significant AI-led financial market dislocation in the 2027–2030 window is material and rising, and current market consensus is, in the author's view, dangerously complacent.

6. References

All references include full URLs verified as of April 2026. Note: some institutional reports require free registration for access.

BloombergNEF (2025). *Power for AI: Easier Said Than Built.* BloombergNEF Energy Transition Research, July 2025. URL: https://about.bnef.com/insights/commodities/power-for-ai-easier-said-than-built/

Deloitte (2025). *Can US Infrastructure Keep Up with the AI Economy?* Deloitte Insights, December 2025. URL: https://www.deloitte.com/us/en/insights/industry/power-and-utilities/data-center-infrastructure-artificial-intelligence.html

Dell'Oro Group (2026). *AI Boom Drives Data Center Capex to $1.7 Trillion by 2030.* Press Release, February 2026. URL: https://www.delloro.com/news/ai-boom-drives-data-center-capex-to-1-7-trillion-by-2030/

Eurelectric (2025). *Power Barometer 2025: AI and Energy.* European Electricity Industry Association, Brussels. URL: https://www.eurelectric.org/in-detail/ai-and-energy/

European Commission (2025). *AI Continent Action Plan.* European Commission, April 2025. URL: https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1013

European Parliament (2025). *Briefing: AI and the Energy Sector.* EPRS BRI(2025)775859. URL: https://www.europarl.europa.eu/RegData/etudes/BRIE/2025/775859/EPRS_BRI(2025)775859_EN.pdf

Gartner (2025). *Electricity Demand for Data Centers to Grow 16% in 2025 and Double by 2030.* Gartner Press Release, November 2025. URL: https://www.gartner.com/en/newsroom/press-releases/2025-11-17-gartner-says-electricity-demand-for-data-centers-to-grow-16-percent-in-2025-and-double-by-2030

Goldman Sachs Research (2025). *AI Energy: The Six P's Framework.* Goldman Sachs Exchanges Research Report, June 2025. URL: https://www.goldmansachs.com/insights/articles/ai-energy-infrastructure

IAEA (2025). *Data Centres, AI and Cryptocurrencies Eye Advanced Nuclear.* IAEA Bulletin, 2025. URL: https://www.iaea.org/bulletin/data-centres-artificial-intelligence-and-cryptocurrencies-eye-advanced-nuclear-to-meet-growing-power-needs

IEA (2025a). *Energy and AI.* IEA Special Report, April 2025. Paris: OECD/IEA. URL: https://www.iea.org/reports/energy-and-ai

IEA (2025b). *Overcoming Energy Constraints for Europe's Data Centre Goals.* IEA Commentary, November 2025. URL: https://www.iea.org/commentaries/overcoming-energy-constraints-is-key-to-delivering-on-europe-s-data-centre-goals

IEA (2025c). *Energy Supply for AI.* IEA Website, 2025. URL: https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai

IPCC (2022). *Climate Change 2022: Mitigation of Climate Change.* WG III, AR6. Cambridge University Press. URL: https://www.ipcc.ch/report/ar6/wg3/

JLL (2026). *2026 Global Data Center Outlook.* Jones Lang LaSalle, March 2026. URL: https://www.jll.com/en-us/insights/market-outlook/data-center-outlook

McKinsey Global Institute (2025). ***The Cost of Compute: A $7 Trillion Race to Scale Data Centers.*** **McKinsey & Company, May 2025. URL: https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers**

NEA/OECD (2023). ***Small Modular Reactors: Challenges and Opportunities.*** **OECD Publishing, Paris. URL: https://www.oecd-nea.org/jcms/pl_64297/small-modular-reactors-challenges-and-opportunities**

NRC (2025). ***Standard Design Approval: NuScale US460 SMR.*** **NRC Publication, May 2025. URL: https://www.nrc.gov/reactors/new-reactors/smrs/nuscale.html**

Oxford Institute for Energy Studies (2026). ***The China Data Centre Advantage.*** **OIES Energy Comment, February 2026. URL: https://www.oxfordenergy.org/wpcms/wp-content/uploads/2026/02/Comment-The-China-data-centre-advantage.pdf**

Pew Research Center (2025). ***Energy Use at US Data Centers Amid the AI Boom.*** **October 2025. URL: https://www.pewresearch.org/short-reads/2025/10/24/what-we-know-about-energy-use-at-us-data-centers-amid-the-ai-boom/**

S&P Global (2025). ***Energy Utility Capex to Top $1 Trillion 2025–2029.*** **S&P Global Market Intelligence, April 2025. URL: https://www.spglobal.com/marketintelligence/en/mi/research-analysis/energy-utility-capex-predicted-to-top-1-trillion-from-2025-through-2029.html**

US DoE (2025). ***GAIN Programme: 2025 Annual Report.*** **US DoE Office of Nuclear Energy, 2025. URL: https://gain.inl.gov/**

Wood Mackenzie (2025). ***Powering China's Data Centres.*** **Wood Mackenzie Energy Research, July 2025. URL: https://www.woodmac.com/blogs/the-edge/powering-chinas-data-centres/**