CSAI Foundation | Cloud Security Alliance

When AI Power Concentration Becomes Systemic Risk

2026-07-17

AI-assisted Rapid Research

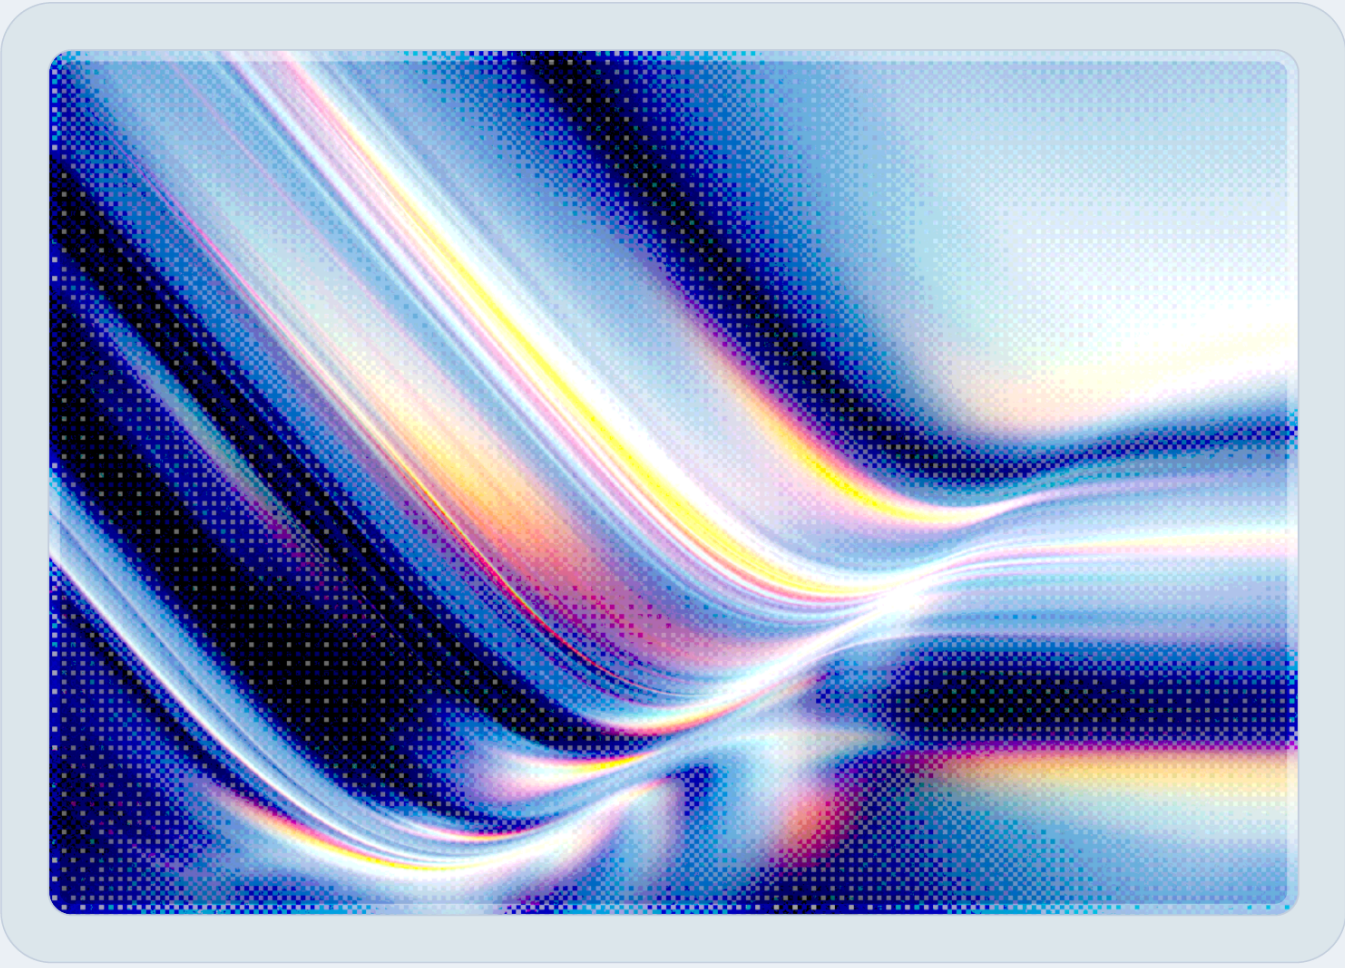

CSAI

CSA cloud security alliance®

Key Takeaways

An essay by security technologist Bruce Schneier and co-author Nathan E. Sanders argues that the public fight against individual AI data center projects, while locally legitimate, distracts from a larger structural problem: a small number of AI companies and their hyperscaler partners are accumulating financial, computational, and political power at a pace with few historical precedents [1]. For security and risk professionals, this argument is not primarily a political or environmental claim; CSA's assessment is that it describes concentration risk building across three interlocking layers—financial markets, physical infrastructure, and governance influence—each of which shows measurable stress, as detailed below. The Bank for International Settlements has formally flagged the AI infrastructure buildout as a threat to global financial stability, citing circular financing arrangements and off-balance-sheet debt structures similar in structure to financing patterns that predated the 2008 financial crisis [2][3][8]. Insurers, meanwhile, are struggling to model accumulation risk across data center clusters that concentrate tens of billions of dollars in replacement value within twenty-mile radiuses of one another [4]. And AI companies themselves have begun spending tens of millions of dollars to shape the state and federal regulatory environment that will govern them, most visibly in New York's 12th congressional district primary [5][6][11]. Enterprises that have built critical workflows atop a handful of AI vendors and the cloud infrastructure beneath them should treat this convergence as a single, compounding risk category rather than three unrelated news stories.

Background

The scale of AI infrastructure investment in 2026 is without precedent in the hyperscalers' own capital spending history. The five largest U.S. hyperscalers—Amazon, Alphabet, Microsoft, Meta, and Oracle—are on pace to spend more than a trillion dollars combined on AI-related capital expenditure across 2025 and 2026, with roughly three-quarters of that spending directed at AI-specific infrastructure such as data centers, chips, and power generation [3][7]. Schneier and Sanders frame this buildout not as an isolated infrastructure story but as the leading edge of a broader ambition: AI firms are explicitly targeting the value created by entire industries, from enterprise software and creative design to law, medicine, and education, sectors that collectively dwarf the data center construction spend itself [1]. CSA's view is that local opposition to individual data centers—citing water use, electricity costs, and noise

—reflects real and immediate community concerns, but the essay's central claim is that this opposition does little to constrain the underlying dynamic of wealth and decision-making authority consolidating in a small number of firms.

That consolidation is now visible in financial markets in ways that go beyond rhetoric. The BIS's 2026 Annual Economic Report named AI infrastructure sustainability as one of four pressure points threatening the global economy, alongside inflation, strained public finances, and broader financial vulnerabilities [8]. The report describes "circular financing," in which chipmakers and hyperscalers take equity stakes in AI labs or neocloud providers that, in turn, commit to multi-year purchases of chips or computing power from those same investors, and warns that many data center construction deals are structured through special purpose vehicles or joint ventures whose lease terms are poorly disclosed, raising the possibility that the same underlying asset is pledged as collateral more than once [2][3]. Private credit funds have quadrupled their lending to AI and IT companies over the past five years, and AI now accounts for roughly 15 percent of their loan books, a pattern CSA reads as evidence that underwriting standards have not kept pace with the added concentration [3]. None of this is a prediction that a crash is imminent, but it suggests that the AI buildout's financing structure shares structural features—opacity and interconnection—with patterns central banks have flagged as systemically dangerous before.

Physical infrastructure concentration compounds the financial exposure. Insurance industry analysis finds that global data center premiums are projected to more than double, from $10.6 billion to $24.2 billion by 2030, as underwriters try to price risk they increasingly cannot cleanly separate across policies [4]. Individual AI-optimized facilities can cost on the order of $20 billion to build and equip, and developers frequently site multiple such facilities within a twenty-mile radius of one another to share power and network infrastructure, as seen in build-outs around Abilene, Texas [4]. Over a quarter of U.S. data center capacity sits in regions prone to large hail, and roughly 40 percent sits in significant tornado zones, meaning a single severe weather event could generate claims against several billion-dollar facilities simultaneously [4]. AI-specific hardware compounds the hazard profile: AI servers draw 100 or more kilowatts compared to 5 to 15 kilowatts for traditional servers, an increasing share of new capacity includes on-site power generation with associated fire and explosion risk, and lithium-ion battery backup units integrated into server racks are a documented and growing ignition source, with fire now driving 42 percent of data center insurance losses despite accounting for only 11 percent of loss events [4].

The third layer—political and regulatory influence—determines whether the first two are ever meaningfully constrained. AI-aligned super PACs spent more than $50 million on 2026 U.S. elections, split between a "pro-innovation" bloc anchored by OpenAI co-founder Greg Brockman's Leading the Future PAC and a "pro-safety" bloc backed by Anthropic-linked committees, including roughly $19 million from Public First on the Bores race in New York's 12th congressional district and $16.6 million from the Jobs and Democracy PAC and Defending Our Values across congressional races more broadly [5][6].

The two blocs' spending competed most visibly in what press coverage described as a proxy war in New York's 12th congressional district, where outside AI-industry spending exceeded $27 million, making it the second-most-expensive U.S. House primary on record, largely over the candidacy of Alex Bores, a former Palantir engineer and author of New York's RAISE Act [6][10][11]. That an industry generating this much capital expenditure is simultaneously willing to spend tens of millions of dollars shaping the individual legislative races that will decide its regulatory environment is, in Schneier and Sanders's framing, the clearest evidence that data center siting fights are a symptom rather than the disease [1].

Security Analysis

Security and risk teams should recognize this convergence as analogous to, though not identical in mechanism to, the concentration-risk pattern CSA has applied to agentic AI vendor consolidation. Both involve a small number of providers accumulating disproportionate control over infrastructure that a larger population of organizations depends on operationally, but the underlying transmission channels here—credit markets, catastrophe risk, and political influence—are distinct from vendor lock-in and warrant separate assessment methods. CSA's own analysis of the U.S. Department of Defense's rapid, large-scale agentic AI deployment documented how quickly organizations can move from adoption to functional dependency to negotiation asymmetry to disruption, and found that only 21 percent of organizations report a mature agentic AI governance model even as nearly half report at least one business function dependent on a single AI vendor. The AI infrastructure concentration described by Schneier and Sanders operates one layer beneath that vendor relationship: it is not just that an enterprise depends on one model provider, but that the model provider itself depends on financing structures, power grids, and hardware supply chains that are themselves concentrated and increasingly systemically fragile.

This likely creates a compounding dependency chain that many enterprise risk registers do not yet capture, extending the governance gap CSA has already documented at the vendor level. A regulatory shock, a credit event in the private lending markets feeding AI infrastructure, or a physical loss event affecting a geographically clustered set of data centers could each independently degrade the availability, pricing, or terms of service an enterprise's AI vendor offers—separately from and in addition to the vendor-level disruption risks (outages, contract disputes, export control designations) that CSA has previously documented. The U.S. Department of War's March 2026 supply chain risk designation against Anthropic and the ten-hour global ChatGPT outage of June 2025 already demonstrated that vendor-level disruption is not theoretical [9][12]. The BIS and insurance industry findings indicate that infrastructure-level disruption, correlated across multiple vendors that share the same hyperscaler data centers, financing vehicles, and power markets, is a distinct and additional exposure that CSA has not seen reflected in the enterprise AI governance programs reviewed to date. An organization that has

diversified across two or three model providers may still be fully exposed if those providers lease capacity from the same handful of hyperscalers, draw power from the same regional grid, or depend on the same private credit facilities.

The political dimension introduces a further, less quantifiable risk. CSA's assessment is that regulatory environments shaped heavily by the spending of the firms being regulated tend to be less predictable for compliance planning, since the outcome depends on contested electoral and legislative fights rather than a stable rulemaking process. Organizations building long-horizon AI governance programs around an assumption of stable, industry-informed regulation should recognize that state-level frameworks like New York's RAISE Act are themselves contested terrain, with well-funded campaigns actively working to defeat or reshape the officials who author them [6][10]. This does not mean such frameworks will fail, but it means enterprises should treat the current regulatory landscape as more volatile than it may appear, and should not build compliance programs that depend on a single anticipated regulatory outcome.

Recommendations

Immediate Actions

Enterprises should extend existing AI vendor risk assessments to explicitly map the infrastructure dependencies beneath each AI vendor relationship: which hyperscaler(s) host the vendor's primary compute, whether that capacity is concentrated in a specific geographic cluster, and whether the vendor or its infrastructure partners have disclosed exposure to the financing structures the BIS has flagged. Security and risk teams should also review whether business continuity and incident response plans account for infrastructure-level disruption (regional weather events, power constraints, credit events affecting an infrastructure partner) as a distinct scenario from vendor-level disruption (outages, contract termination, pricing changes), since the two require different mitigations and different monitoring signals.

Short-Term Mitigations

Organizations should incorporate infrastructure concentration indicators into ongoing AI vendor due diligence, treating geographic clustering of a vendor's compute capacity and reliance on off-balance-sheet financing structures as material risk factors alongside the operational and security criteria CSA has previously recommended. Where multiple AI vendors are used for genuine architectural diversification rather than contractual diversification alone, teams should verify that the vendors do not share underlying infrastructure providers or financing counterparties to the degree that a single systemic event

could degrade all of them simultaneously. Enterprises should also monitor public disclosures from AI vendors and their infrastructure partners regarding capital structure, lease terms, and power supply arrangements as a standing component of vendor risk review, not a one-time assessment.

Strategic Considerations

Boards and executive risk committees should recognize AI infrastructure concentration as a category of systemic risk that sits above individual vendor risk, comparable in structure to the interconnected counterparty risk regulators have long monitored in banking. This argues for treating AI vendor and infrastructure diversification as a resilience investment rather than a procurement cost, consistent with CSA's prior guidance on architectural (not merely contractual) diversification and interoperability standards that preserve model-backend substitutability. Organizations with substantial AI-dependent operations should also track the political and regulatory environment shaping AI governance as a risk input in its own right, since a regulatory framework can be contested and reshaped by the same concentration of capital that created the infrastructure dependency in the first place.

CSA Resource Alignment

This analysis extends CSA's existing concentration-risk research rather than introducing a new framework. CSA's prior analysis of the U.S. Department of Defense's rapid, large-scale agentic AI deployment documented how such adoption produces an "enterprise pattern" of dependency, negotiation asymmetry, and disruption discovery, and mapped operational and strategic concentration risk to CSA's governance frameworks. The present analysis adds a layer beneath that vendor-level pattern: the financing, physical infrastructure, and political economy underneath the vendors themselves are subject to the same concentration dynamic, and enterprises that have only assessed vendor-level dependency may be missing a material portion of their exposure. The distinction between operational concentration (dependency on a single platform) and strategic concentration (a thin vendor architecture for mission-critical functions) established in that analysis applies directly to infrastructure: enterprises should ask not only which AI vendors they depend on, but which hyperscalers, financing structures, and power markets those vendors themselves depend on.

The AI Controls Matrix (AICM) v1.1 remains the appropriate control framework for translating this risk into governance action, particularly its supply chain security and business continuity domains, which require organizations to assess and document dependency on third-party AI infrastructure and to maintain continuity provisions comparable to those used for other critical infrastructure vendors. Organizations

building or refreshing AI vendor risk assessment programs should use AICM's supply chain and continuity control objectives as the baseline against which infrastructure concentration indicators, such as those described above, are formally tracked and escalated.

References

[1] Schneier, Bruce, and Nathan E. Sanders. "AI Data Centers and the Concentration of Wealth." Schneier on Security, July 2026.

[2] Crypto Briefing. "BIS Warns AI Investment Boom Could End in a Bust." Crypto Briefing, June 2026.

[3] Fortune. "The Central Bank of Central Banks Just Released Its Flagship Annual Report — And It Sees a $1 Trillion AI Investment Boom Headed for a Reckoning." Fortune, June 2026.

[4] Risk & Insurance. "Data Centers Powering AI Create Unprecedented Risk Accumulation Challenges for Insurers." Risk & Insurance, 2026.

[5] NPR. "Groups Tied to OpenAI and Anthropic Are Spending Big on the Midterms." NPR, June 2026.

[6] Fortune. "Anthropic and OpenAI's $50 Million Election Battlefront Has No Winners, and NY-12 Is One Example Why." Fortune, June 2026.

[7] Ropes & Gray LLP. "Data Center Investment in 2026: AI Demand, Power Constraints, and Private Equity Trends." Ropes & Gray, 2026.

[8] Bank for International Settlements. "Global Economic Pressure Points Call for Policy Discipline: BIS Annual Economic Report 2026." BIS Press Release, June 2026.

[9] Mayer Brown. "Anthropic Supply Chain Risk Designation Takes Effect — Latest Developments and Next Steps for Government Contractors." Mayer Brown Insights, March 2026.

[10] ThePrint. "Who Is Alex Bores? AI Firms Are Spending Millions to Defeat Him in New York Congressional Race." ThePrint, 2026.

[11] Fox News. "New York Congressional Primary Becomes Second Most Expensive House Race Ever Amid AI Fight." Fox News, June 2026.

[12] Tom's Guide. "ChatGPT Is Back Following Global Outage — Here's What Happened." Tom's Guide, June 2025.

