CSAI Foundation | Cloud Security Alliance

AI Capital Concentration and the Regulatory Capture Risk

2026-07-13

AI-assisted Rapid Research

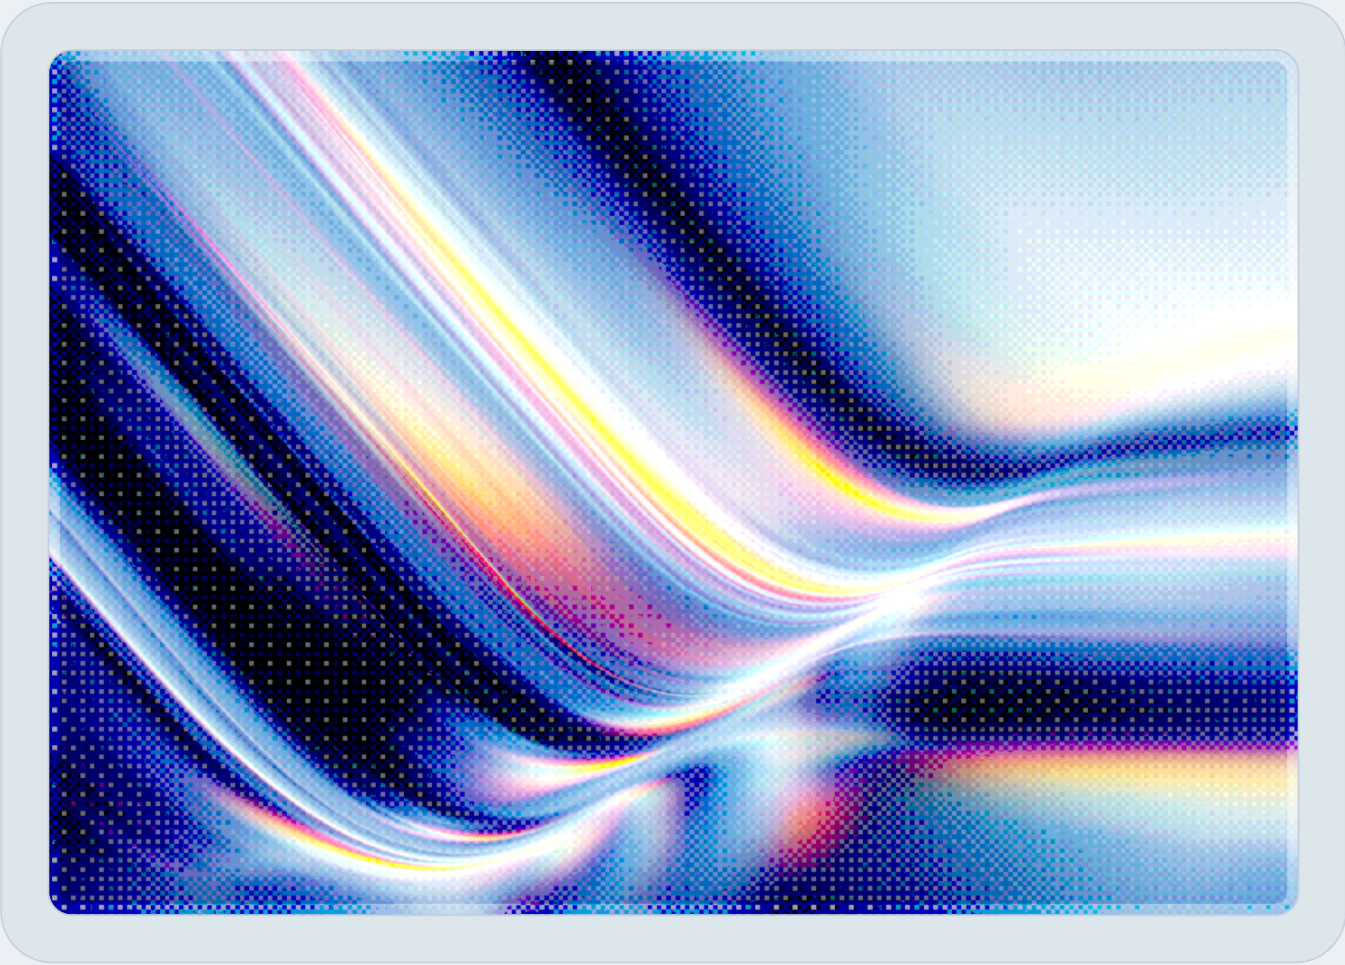

CSAI

CSA cloud security alliance®

Key Takeaways

- A July 9, 2026 essay by security technologist Bruce Schneier and data scientist Nathan E. Sanders argues that public opposition to AI data centers, while legitimate, distracts from a larger and more durable risk: a small number of AI companies accumulating enough capital and market power to shape the regulatory and political environment governing their own industry [1].
- That concentration is measurable. KPMG's Q1 2026 venture capital data shows OpenAI, Anthropic, and xAI alone accounted for roughly $172.6 billion of the $206 billion raised in ten mega-deals that quarter — about 84% of all mega-deal funding and just over half of the quarter's total global VC investment of $330.9 billion [2]. The same five hyperscalers building the infrastructure beneath these models — Microsoft, Alphabet, Meta, Oracle, and Amazon — have projected a combined $720 billion in 2026 capital expenditure [3].
- Peer-reviewed research accepted at the 2026 ACM Conference on Fairness, Accountability, and Transparency identifies 27 distinct types of mechanism — spanning lobbying, revolving-door hiring, standard-setting participation, and funding of the academic research that evaluates their own products — by which AI firms are already shaping the rules meant to govern them, and catalogues 249 individual instances of those mechanisms across the news coverage the researchers analyzed [4]. Independent reporting has compared these tactics directly to the regulatory playbooks historically used by the tobacco, pharmaceutical, and oil industries [5].
- That influence is now visible in concrete outcomes: OpenAI's federal lobbying grew roughly sevenfold year over year through Q1 2026, with Anthropic's spending surpassing OpenAI's for the first time in the same quarter [6]; AI-industry-linked super PACs spent more than $27 million contesting a single New York congressional primary over AI safety legislation [7]; and a December 2025 executive order established a Department of Justice task force whose sole purpose is challenging state AI laws in federal court [8].
- For security and risk leaders, the relevant conclusion is not that regulatory capture guarantees favorable rules for AI vendors — the New York race shows industry spending can fail to produce its intended outcome [7] — but that the AI industry's economic concentration gives it durable, well-resourced influence over the same policy processes that would otherwise set baseline security, safety-testing, and incident-disclosure obligations for the AI systems enterprises depend on.

Background

KPMG's Venture Pulse data for the first quarter of 2026 recorded $330.9 billion in global venture capital investment, a new quarterly record driven almost entirely by AI: ten rounds of $2 billion or more contributed more than $206 billion of that total, and the seven largest of those rounds — led by OpenAI's $122 billion raise, Anthropic's $30.6 billion raise, and xAI's $20 billion raise — together accounted for roughly $198.5 billion, or about 96% of the mega-deal total [2]. On the infrastructure side of the same buildout, the five U.S. hyperscalers most exposed to AI compute demand have told investors they will collectively spend approximately $720 billion on capital expenditure in 2026, a figure substantially larger than prior technology infrastructure buildouts and one that Schneier and Sanders cite as evidence the industry is racing to lock in physical and economic advantages before markets or regulators can catch up [1][3]. This is the backdrop against which Schneier and Sanders wrote their July 2026 essay: they argue that fights over individual data centers, while grounded in genuine local harms such as strained power grids and minimal local employment, are strategically convenient for AI companies because they channel political energy toward site-specific infrastructure battles rather than toward the underlying question of how much economic and political power a handful of firms are accumulating [1]. Their essay specifically points to the industry's stated ambition to capture value currently generated across teaching, medicine, legal services, and creative professions as the more consequential story, and one that opposition to individual facilities does little to address [1].

That underlying concern — that concentrated capital converts into concentrated political influence — has independent empirical support. A paper by researchers at the University of Edinburgh, Trinity College Dublin, TU Delft, and Carnegie Mellon University, accepted at the 2026 ACM Conference on Fairness, Accountability, and Transparency, analyzed news coverage of AI policy debates across the EU AI Act negotiations and AI summits held in the UK, South Korea, and France between 2023 and 2025 [4]. The researchers identified 27 distinct types of capture mechanism, organized into five categories, and catalogued 249 individual instances of those mechanisms across the coverage they reviewed: direct influence on policy through lobbying and political contributions; conflicted involvement through revolving-door appointments and direct participation in rulemaking; market influence through standard-setting bodies and pressure on smaller firms to oppose regulation; elusion of law through jurisdiction shopping and contested interpretations of antitrust, privacy, and copyright statutes; and epistemic influence through corporate funding of the academic research and public discourse that shapes how AI risk is understood [4]. The most frequent category by far was epistemic and discourse influence, and the dominant narrative frames the researchers identified — that "regulation stifles innovation," that compliance amounts to unnecessary "red tape," and that national competitiveness requires deregulation — appeared disproportionately in coverage most closely tied to industry sources [4]. Independent technology press coverage of the same research drew an explicit comparison to the historical playbooks

of the tobacco, pharmaceutical, and fossil fuel industries, each of which used similar combinations of funded research, narrative framing, and personnel movement between regulators and industry to delay or weaken oversight for decades [5].

The political dimension of this dynamic has moved from academic description to observable spending in 2026. Anthropic's federal lobbying expenditure overtook OpenAI's for the first time in the first quarter of 2026, at $1.6 million to $1.5 million, with OpenAI's spending having grown roughly sevenfold from about $210,000 in the same quarter of 2025 [6]. Both companies have also moved substantial money into electoral politics: political groups linked to the two firms spent a combined $27 million on a single New York Democratic congressional primary in June 2026, split between roughly $19 million from an Anthropic-linked group supporting a state legislator who had co-sponsored New York's AI safety-incident-reporting law and roughly $8 million from an OpenAI-linked group opposing him [7]. Neither side's preferred candidate ultimately won the primary, a result commentary on the race described as evidence that heavy AI-industry spending does not reliably convert into electoral control, even as it demonstrates the industry's willingness to deploy tens of millions of dollars around a single House seat [7]. At the federal executive level, a December 11, 2025 executive order titled "Ensuring a National Policy Framework for Artificial Intelligence" directed the Department of Justice to stand up an AI Litigation Task Force dedicated to challenging state AI laws in federal court, instructed the FCC and FTC to pursue preemptive federal standards, and conditioned a portion of Broadband Equity, Access, and Deployment program funding on states repealing AI rules the administration deems burdensome — while explicitly carving out state child-safety, data-center, and government-procurement rules from preemption [8].

Security Analysis

For CSA's audience, the security-relevant question is not whether AI companies are engaging in ordinary industry advocacy — all regulated industries lobby — but what happens to the baseline security expectations that enterprises implicitly rely on when the rulemaking process itself is shaped disproportionately by the firms whose products would be regulated. State AI safety laws that have advanced in the absence of a comprehensive federal framework, such as incident-reporting and safety-testing requirements, function for many enterprises as a de facto minimum bar for vendor accountability, similar to the role state breach-notification laws have played in the absence of comprehensive federal legislation. A litigation and preemption strategy explicitly designed to challenge those state laws, paired with the epistemic-influence patterns documented in the FAccT research — funding studies, shaping conference agendas, and placing personnel in standard-setting bodies — creates a real possibility that the regulatory floor enterprises are counting on for vendor security and safety assurances erodes or fails to solidify at the pace the underlying technology is being deployed [4][8]. This is not a hypothetical concern layered on top of unrelated risk; it compounds a related concentration risk at the infrastructure

layer, where a small number of AI model and cloud providers hold outsized structural influence over enterprise AI dependencies, creating correlated failure modes if any single vendor experiences disruption, a change in terms, or a regulatory designation that limits its availability.

The mechanism connecting capital concentration to security posture runs through two channels. First, market concentration itself is a security variable: when the overwhelming share of frontier AI capability and infrastructure investment concentrates in three or four firms — as the KPMG venture data and hyperscaler capex figures both show — enterprises inherit vendor concentration risk regardless of how the regulatory environment evolves, because architectural alternatives narrow as capital consolidates around fewer providers [2][3]. Second, and more directly tied to the events described above, political and regulatory capture risk affects which security obligations become binding in the first place. The FAccT taxonomy's "elusion of law" category — contested interpretations of existing statutes, jurisdiction shopping, and disregard for antitrust or data-protection requirements — describes behavior that, if it succeeds, would leave enterprises relying on AI vendors' voluntary security commitments rather than externally enforceable ones [4]. The New York primary result is a useful corrective against overstating this risk: heavy spending did not secure the outcome either AI company sought, suggesting capture is neither automatic nor guaranteed [7]. But the scale of spending, the growth rate of federal lobbying budgets, and the explicit design of the December 2025 executive order to challenge state authority collectively indicate that at least the largest AI firms are treating influence over their own regulatory environment as a first-order strategic priority, not an afterthought, and enterprises that assume regulation will keep pace with deployment risk being caught by a governance gap shaped in part by both industry lobbying and executive action.

Recommendations

Immediate Actions

Security and compliance teams should inventory which of their AI governance controls currently rely on regulatory mandates — such as state incident-reporting or model-testing requirements — versus vendor-provided voluntary commitments, and flag the former as subject to change given active federal litigation and preemption efforts targeting state AI statutes [8]. Enterprises operating in states with AI safety-incident-reporting requirements should continue complying with those obligations as currently written rather than anticipating preemption, since the December 2025 executive order lacks independent legal force absent a statute or judicial ruling and existing state laws remain binding until successfully challenged [8].

Short-Term Mitigations

Organizations should treat AI vendor concentration and AI regulatory capture as related but distinct entries in their third-party risk register: the former assessed through vendor dependency and switching-cost analysis, and the latter through monitoring of pending state and federal AI legislation, executive actions, and litigation that could alter vendor security or disclosure obligations with limited notice. Risk and legal teams should also track which specific security-relevant provisions — incident reporting, third-party audit rights, safety-testing transparency — exist in current state law and would be lost if federal preemption succeeds, so that any gap can be addressed through contractual terms with AI vendors rather than discovered after the fact.

Strategic Considerations

Given that the FAccT research identifies epistemic and discourse influence — funded research, narrative framing, conference sponsorship — as the most prevalent capture mechanism, enterprises and industry bodies engaged in AI security standard-setting should prioritize participation from organizations without direct commercial stakes in frontier model development, and should scrutinize the provenance of safety and risk research cited in policy debates [4]. Because political spending on AI regulation has not reliably produced the outcomes its sponsors sought, as the New York race illustrates, enterprises and civil-society voices retain meaningful ability to shape outcomes if they engage with the same intensity industry has brought to lobbying and electoral spending [7]. Enterprises should not treat the current absence of comprehensive federal AI regulation as a stable baseline; they should instead build internal governance, vendor risk management, and control frameworks robust enough to hold regardless of how the external regulatory floor moves in either direction over the next 12 to 24 months.

CSA Resource Alignment

CSA's "AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects" provides the organizational governance foundation most directly relevant to this note's core recommendation: that enterprises build internal AI risk management and compliance programs that do not depend on the trajectory of external regulation remaining stable, since that guide's framework for board-level ownership and cultural accountability of AI risk is designed to function independently of whichever regulatory floor is ultimately in force [9]. The AI Controls Matrix (AICM) v1.1 supplies the control-level complement to that governance layer; its supply chain and third-party risk domains give organizations a structured way to formalize the vendor concentration and dependency analysis this note recommends, translating the capital-concentration dynamic described here into concrete controls

around vendor diversification, contractual continuity provisions, and audit rights that do not rely on regulatory mandates to be enforceable [10]. Enterprises applying either framework should treat the erosion or delay of state-level AI safety statutes discussed above as a reason to specify security and disclosure obligations directly in vendor agreements and internal controls, rather than as a gap to be filled by future regulation.

References

[1] Bruce Schneier and Nathan E. Sanders. "The Fight Against AI Data Centers Is Important — but It's Just a Starting Point." The Guardian / Schneier on Security, July 9, 2026.

[2] KPMG. "Global VC Investment Surges to Record $330.9 Billion in Q1'26 on Back of AI Megadeals." KPMG Private Enterprise Venture Pulse, April 2026.

[3] The Motley Fool. "The $720 Billion Capex Trap: 2 Artificial Intelligence (AI) Hyperscalers Spending on Growth While the Rest Spend on Maintenance." The Motley Fool, April 25, 2026.

[4] Abeba Birhane, Riccardo Angius, William Agnew, Harshvardhan J. Pandit, Bhaskar Mitra, Roel Dobbe, and Zeerak Talat. "Big AI's Regulatory Capture: Mapping Industry Interference and Government Complicity." Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26), May 2026.

[5] The Register. "'Big AI' Is Subverting Regulations Just Like Tobacco and Oil Firms." The Register, May 18, 2026.

[6] The AI Lobby. "Anthropic Outspends OpenAI: The AI Lobbying Arms Race Heats Up (Q1 2026)." The AI Lobby, 2026.

[7] Fortune. "Anthropic and OpenAI Waged a $27 Million Proxy War in a Manhattan Congressional Race. The Winner Told Them Both to Get Lost.." Fortune, June 26, 2026.

[8] Gibson Dunn. "President Trump's Latest Executive Order on AI Seeks to Preempt State Laws." Gibson Dunn Client Alert, December 2025.

[9] Cloud Security Alliance. "AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects." Cloud Security Alliance, 2024.

[10] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.