Governing the Ungoverned: U.S. AI Security Policy Fragmentation

Navigating the Trump Cyber Strategy, NIST AI Agent Standards, and CISA's Capacity Crisis

Unofficial AI-assisted Research

Cloud Security Alliance AI Safety Initiative

2026-03-20

Key Takeaways

- The Trump administration's March 2026 national cyber strategy embraces voluntary industry frameworks and rejects mandatory compliance, creating a period in which enforceable requirements are replaced by voluntary measures at a time when AI-powered threats are maturing rapidly.
- NIST's Center for AI Standards and Innovation (CAISI) launched a formal AI Agent Standards Initiative in February 2026, representing the federal government's most direct technical engagement with agentic AI security risks under the current administration—but the initiative lacks enforcement authority.
- CISA has lost approximately one-third of its workforce since early 2025 [1], with critical programs in election security, national risk management, and cyber defense training eliminated or severely curtailed. The agency's capacity to coordinate AI security incident response is materially degraded.
- No federal statute governs AI security. The Biden-era OMB guidance (M-24-10) requiring agencies to adopt NIST AI Risk Management Framework principles has an uncertain enforcement status under the current administration, and no successor legislation has passed Congress.
- Enterprises and security practitioners face a fragmented landscape: a deregulatory federal posture, active NIST standards activity, weakened federal incident coordination, and more than 100 state-level AI measures across 38 jurisdictions [2]. Organizations cannot assume federal guidance will provide a coherent baseline.

Background

The U.S. AI security governance landscape has undergone three significant structural changes within fourteen months: the revocation of the Biden administration's primary AI safety framework, a major institutional restructuring at the agency responsible for evaluating frontier AI models, and the release of a replacement national cyber strategy with a fundamentally different regulatory philosophy. These shifts have unfolded while agentic AI systems—the category that poses the most novel and least-understood security risks—have continued their rapid commercial deployment without a corresponding federal governance structure.

The inflection point was January 20, 2025, when President Trump signed Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence," revoking the Biden administration's EO 14110 [3]. Biden's order had required mandatory red-teaming for high-risk AI models, enhanced cybersecurity protocols for AI in critical infrastructure, and detailed interagency coordination mechanisms for AI-related national security threats. Its revocation was immediate and left no transition period. Agencies that had built compliance processes around EO 14110's requirements found those processes suddenly unmoored from a governing mandate.

The organizational consequences followed quickly. The U.S. AI Safety Institute (AISI), established under the Commerce Department to evaluate frontier AI models and develop safety standards, was renamed the Center for AI Standards and Innovation (CAISI) in June 2025 [4]. Commerce Secretary Howard Lutnick framed the change explicitly: CAISI would focus on "secure innovation" rather than safety oversight, narrowing its mission from broad societal risk mitigation to national security threats, cybersecurity, biosecurity, and countering foreign adversarial influence. The word "safety" was removed from the organization's name—a change Lutnick framed as a mission-scope adjustment, though critics have interpreted it as a deliberate signal about the administration's priorities.

In March 2026, the Trump administration published its first formal national cyber strategy—a seven-page document articulating six pillars: shaping adversary behavior through offensive operations, promoting what it characterizes as "common sense regulation," modernizing federal networks, securing critical infrastructure, maintaining technological superiority, and building cyber talent [5]. Its most significant departure from prior strategies is its explicit rejection of mandatory compliance frameworks. Where the Biden-era National Cybersecurity Strategy called for minimum security requirements in critical infrastructure sectors, the Trump strategy frames compliance-driven approaches as "a costly checklist" and signals reliance on public-private collaboration and voluntary frameworks instead.

This policy pivot lands against a backdrop of unprecedented reductions at the agency most responsible for coordinating domestic cybersecurity. CISA began 2025 with approximately 3,300 to 3,400 employees; by early 2026, that number had fallen to approximately 2,300 to 2,400, a reduction of roughly 29 to 30% [1][14]. DOGE-driven cuts in March 2025 eliminated the agency's red team—more than 100 cybersecurity specialists whose contract was cancelled outright [6]. Subsequent reductions through buyouts, layoffs, and early retirements significantly defunded CISA's Elections Infrastructure Information Sharing and Analysis Center (defunded as of March 2025), its Stakeholder Engagement Division, and its National Risk Management Center—all of which sustained substantial funding reductions [7]. The agency has operated without a Senate-confirmed director since early 2026, with multiple holds on the nomination of Sean Plankey blocking Senate confirmation and Madhu Gottumukkala, serving as acting Deputy Director, filling the leadership role.

Security Analysis

Regulatory Fragmentation

The net effect of these concurrent changes—policy revocation, institutional restructuring, and workforce reduction—is a regulatory environment whose dimensions are not yet fully visible but whose contours are becoming clear. The United States currently has no binding federal statute governing AI security. The Biden administration's M-24-10 OMB guidance, which required federal agencies to implement NIST AI RMF principles by December 2024, technically remains on the books but lacks clear enforcement authority under the current administration. The TAKE IT DOWN Act, the only AI-specific federal legislation enacted in 2025, addresses nonconsensual intimate images and has no bearing on AI security [8]. More than 1,000 AI-related bills were introduced across U.S. states and territories during 2025, and approximately 38 jurisdictions enacted roughly 100 AI measures [2]—but these vary widely in scope, enforcement mechanisms, and applicability to security-specific concerns. The December 2025 EO 14365 directed the Department of Justice's newly formed AI Litigation Task Force to challenge state AI laws perceived as conflicting with federal interests, initiating what legal analysts describe as a coming wave of preemption litigation [9].

For enterprises deploying AI systems in critical infrastructure, financial services, healthcare, or federal contracting contexts, this fragmentation creates material compliance risk. Security teams cannot identify a single authoritative federal framework against which to audit AI deployments. The risk of investing in compliance with a state-level requirement that is subsequently invalidated through federal preemption is real, as is the inverse risk of deferring to a permissive federal posture while a state in which the enterprise operates imposes stricter obligations.

NIST's AI Agent Standards Initiative: Substantive but Voluntary

The most technically substantive federal response to agentic AI security risks in this period has come from NIST and CAISI. In January 2026, CAISI published a Request for Information on security considerations for AI agent systems in the Federal Register, soliciting input on threat models, autonomous task execution vulnerabilities, human supervision mechanisms, and access control approaches [10]. The RFI attracted substantial comment and informed a formal initiative announced February 17, 2026.

The AI Agent Standards Initiative addresses three areas: facilitating U.S. leadership in international AI standards bodies (explicitly framed as countering China's growing influence in ISO/IEC forums [11]), supporting open-source protocol development for agent interoperability, and advancing research in AI

agent security and identity. Concurrently, the National Cybersecurity Center of Excellence (NCCoE) released a concept paper on AI agent identity and authorization, exploring OAuth 2.0-based authorization mechanisms, agent identity distinct from human user identities, and audit logging requirements for non-human entities [12]. NIST IR 8596, a preliminary draft Cyber AI Profile integrating NIST Cybersecurity Framework 2.0 with AI-specific risks, attracted more than 6,500 members to its community of interest when published in December 2025 [13].

These are meaningful technical contributions. The NCCoE's framing of agent identity as a distinct authorization problem—separate from the human users who delegate tasks to agents—directly addresses one of the most acute security gaps in current agentic deployments, where agents often inherit ambient credentials without explicit scoping or time-bounding. The Cyber AI Profile's three-focus structure (securing AI systems, AI-enabled cyber defense, and defending against AI-enabled attacks) provides a coherent analytical frame. However, CAISI and NIST operate as standards-setting bodies, not regulators. Their outputs are voluntary guidance and aspirational frameworks. Nothing in the AI Agent Standards Initiative creates an enforceable obligation for organizations deploying AI agents, and the administration's broader posture explicitly discourages mandatory requirements.

This creates an important structural tension: the entity doing the most rigorous technical work on AI agent security (NIST/CAISI) lacks the authority to require its own standards' adoption, while the entity with the broadest coordinating authority for domestic cybersecurity (CISA) has seen its operational capacity significantly reduced.

The CISA Capacity Problem

The practical security implications of CISA's workforce and budget reductions extend beyond organizational metrics. CISA's National Risk Management Center, which coordinated sector-specific risk assessments and information sharing between federal agencies and critical infrastructure operators, has seen its operational capacity materially diminished. The agency's ability to conduct joint threat assessments, facilitate cross-sector information sharing on AI-specific attack vectors, or respond rapidly to novel AI security incidents is materially reduced relative to 2024. The elimination of CISA's red team removes a capability specifically designed to test the resilience of federal systems and critical infrastructure—at a moment when adversary use of AI-powered techniques has been increasingly documented in threat reporting.

The leadership vacancy compounds these structural issues. A Senate-confirmed CISA director provides the institutional authority needed to engage at the interagency level, advocate for resources, and set enforcement priorities. Acting leadership, however capable, lacks that mandate. Multiple holds on Plankey's nomination—the reasons for which have not been fully disclosed publicly—have left CISA unable to resolve this uncertainty for more than twelve months [14].

The combined effect is a degraded national capacity to coordinate response to precisely the kinds of incidents that the Trump Cyber Strategy identifies as priorities: AI-enabled attacks on critical infrastructure, adversary use of AI to accelerate vulnerability discovery, and the compromise of federal networks by state-sponsored actors wielding AI-powered intrusion tools.

International Standards Competition as a Governance Driver

One underappreciated dimension of the current fragmentation is the role of geopolitical competition in shaping U.S. standards strategy. CAISI's explicit framing of the AI Agent Standards Initiative as a counter to Chinese influence in international standards bodies reflects a recognition that governance frameworks are themselves a form of strategic competition [11]. Nations and blocs that establish the foundational technical standards for AI agents—identity, authorization, interoperability, audit logging—will shape the security properties of the global AI infrastructure for years. The U.S. approach, driven through NIST's voluntary standards process and international engagement through CAISI, differs fundamentally from the EU's regulatory approach under the EU AI Act, and from emerging frameworks in the Asia-Pacific region. These divergences create cross-jurisdictional compliance challenges for multinational enterprises and increase the difficulty of incident attribution and response coordination.

Recommendations

Immediate Actions

Enterprises and security teams should not wait for federal policy coherence to materialize before establishing AI security governance practices. Organizations should assess their current AI deployments against NIST's AI Risk Management Framework (AI RMF 1.0) and the preliminary Cyber AI Profile (NIST IR 8596), treating voluntary guidance as a practical baseline in the absence of mandatory requirements. Particular attention should be paid to the NCCoE's emerging guidance on AI agent identity and authorization: any agentic system that executes tasks using credentials, makes API calls to external services, or operates with access to sensitive data should be evaluated for alignment with the principles of distinct agent identity, scoped authorization, time-bounded access, and audit logging.

Organizations with federal contracts or critical infrastructure designations should actively monitor the enforcement status of OMB M-24-10 and track whether agency AI governance requirements survive the current policy transition or are formally rescinded. The uncertainty itself is a risk that should be reflected in compliance planning timelines.

Short-Term Mitigations

Because CISA's capacity to provide sector-specific guidance and coordinate incident response has been reduced, enterprises in critical infrastructure sectors should invest in direct participation with sector-specific Information Sharing and Analysis Centers (ISACs) and other industry-led mechanisms to supplement reduced federal coordination capacity and maintain situational awareness. ISAC participation varies in depth across critical infrastructure sectors, and organizations should assess what their relevant ISAC provides before treating it as a direct substitute for federal coordination. The reduction of CISA's Stakeholder Engagement Division means that the informal channels through which CISA historically communicated threat intelligence to industry have been narrowed; organizations should not assume they will receive the same quality of advance warning for emerging AI-specific threats that was available in prior years.

Legal and compliance teams should map their organization's AI deployments against the patchwork of state AI measures currently in effect—particularly in California, Colorado, and New York, which have indicated through ongoing legislative activity their intention to maintain state-level AI requirements even in the face of federal preemption efforts [9]. The outcome of DOJ's AI Litigation Task Force challenges to state laws is unlikely to be resolved quickly, and operating in multiple states increases exposure to conflicting obligations in the interim period.

Strategic Considerations

The current policy environment, while creating compliance complexity, also presents an opportunity for enterprises to shape the voluntary frameworks that are likely to serve as de facto standards in the absence of mandatory regulation. NIST's public comment processes—including the open comment period on the NCCoE's AI agent identity concept paper, which closes April 2, 2026, as of this publication [12]—provide direct channels for practitioners to influence the technical standards that will govern agentic AI security. Organizations with relevant operational experience in deploying AI agents at scale should engage these processes actively.

At the strategic level, boards and executive leadership should treat U.S. AI security governance fragmentation as a material business risk, not merely a compliance inconvenience. The absence of a coherent federal framework increases the probability of regulatory divergence, enforcement uncertainty, and gaps in national incident coordination that could affect enterprise operations during a major AI-related security event. Risk frameworks and board-level reporting should reflect this elevated uncertainty.

CSA Resource Alignment

The governance fragmentation described in this note is precisely the operational environment that CSA's frameworks were designed to help enterprises navigate. The AI Controls Matrix (AICM)—a superset of the Cloud Controls Matrix (CCM)—provides a vendor-neutral control baseline that enterprises can apply regardless of which regulatory framework ultimately prevails. In the current environment, where federal requirements are in flux and state requirements are proliferating, the AICM's control-level specificity allows organizations to demonstrate security posture against a credible third-party standard even in the absence of binding federal mandate.

CSA's MAESTRO framework for agentic AI threat modeling directly addresses the attack surface categories that NIST's AI Agent Standards Initiative and the NCCoE concept paper are attempting to govern: tool abuse, credential inheritance, unauthorized task escalation, and the absence of meaningful human oversight in autonomous execution chains. MAESTRO's threat hierarchy provides a structured approach to the same problems that the federal government is still in the early stages of standardizing, giving enterprises a head start on implementing defensible controls.

The CSA STAR program provides a mechanism for communicating AI security posture to customers, partners, and regulators—including regulators in jurisdictions that may impose disclosure or audit requirements before a coherent federal framework exists. For enterprises operating across multiple regulatory environments, STAR attestations offer a consistent signal that does not depend on the resolution of federal-state preemption disputes.

CSA's Zero Trust guidance connects directly to the authorization and identity challenges that the NCCoE's AI agent concept paper identifies. The principle that no entity—human or non-human—should receive implicit trust based on network location or credential inheritance is precisely the control philosophy needed to address the ambient-credential risks that agentic systems create. Organizations should treat the NCCoE's emerging agent authorization framework as an extension of their existing Zero Trust implementation rather than a parallel initiative.

Finally, CSA's AI Organizational Responsibilities publications provide governance-level guidance on the accountability structures, roles, and escalation processes needed to manage AI systems in production. As CISA's coordinating capacity diminishes and federal incident response structures become less reliable, the internal governance structures described in these publications take on greater operational importance.

References

[1] Nextgov, "CISA Projected to Lose a Third of Its Workforce Under Trump's 2026 Budget," June 2025. https://www.nextgov.com/cybersecurity/2025/06/cisa-projected-lose-third-its-workforce-under-trumps-2026-budget/405726/

[2] Harvard Safra Center for Ethics, "AI Governance at a Crossroads: America's AI Action Plan and Its Impact on Businesses," November 2025. https://www.ethics.harvard.edu/news/2025/11/ai-governance-crossroads-americas-ai-action-plan-and-its-impact-businesses

[3] White House / Federal Register, "Executive Order 14179: Removing Barriers to American Leadership in Artificial Intelligence," January 23, 2025. https://www.federalregister.gov/documents/2025/01/28/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence

[4] U.S. Department of Commerce, "Statement of Secretary Howard Lutnick: Transforming the U.S. AI Safety Institute," June 3, 2025. https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai

[5] Mayer Brown, "Trump Administration Releases Cyber Strategy for America and Related Executive Order on Combatting Cybercrime," March 2026. https://www.mayerbrown.com/en/insights/publications/2026/03/trump-administration-releases-cyber-strategy-for-america-and-related-executive-order-on-combatting-cybercrime

[6] TechCrunch, "DOGE Axes CISA Red Team Staffers Amid Ongoing Federal Cuts," March 11, 2025. https://techcrunch.com/2025/03/11/doge-axes-cisa-red-team-staffers-amid-ongoing-federal-cuts/

[7] CyberScoop, "House Sets CISA Budget Cut at $135 Million," July 2025. https://cyberscoop.com/cisa-budget-135m-dhs-2026-house-appropriations/

[8] Congress.gov, "TAKE IT DOWN Act (S.146, 119th Congress)," enacted May 19, 2025. https://www.congress.gov/bill/119th-congress/senate-bill/146

[9] Sidley Austin LLP, "Unpacking the December 11, 2025 Executive Order: Ensuring a National Policy Framework for Artificial Intelligence," December 23, 2025. https://datamatters.sidley.com/2025/12/23/unpacking-the-december-11-2025-executive-order-ensuring-a-national-policy-framework-for-artificial-intelligence/

[10] Federal Register, "Request for Information Regarding Security Considerations for Artificial Intelligence Agents," NIST/CAISI, January 8, 2026. https://www.federalregister.gov/documents/2026/01/08/2026-00206/request-for-information-regarding-security-considerations-for-artificial-intelligence-agents

[11] NIST, "Announcing the AI Agent Standards Initiative: Interoperable and Secure," February 17, 2026. https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure

[12] NCCoE/NIST, "Accelerating the Adoption of Software and AI Agent Identity and Authorization," Concept Paper, February 2026. https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf

[13] NIST, "NIST IR 8596 (Initial Public Draft): Cybersecurity Framework Profile for Artificial Intelligence," December 16, 2025. https://csrc.nist.gov/pubs/ir/8596/iprd

[14] Federal News Network, "CISA Director Void Leaves Cyber Agency Embroiled in Uncertainty," January 2026. https://federalnewsnetwork.com/cybersecurity/2026/01/cisa-director-void-leaves-cyber-agency-embroiled-in-uncertainty/

