US Federal AI Security Governance in Crisis: CISA Capacity, Pentagon AI Policy, and the Responsible Scaling Vacuum

US Federal AI Security Governance in Crisis: CISA Capacity, Pentagon AI Policy, and the Responsible Scaling Vacuum

Assessing the Structural Gaps in America's AI Safety Architecture

Unofficial AI-assisted Research

Cloud Security Alliance AI Safety Initiative

2026-03-07

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Executive Summary

The federal architecture for AI security governance in the United States is undergoing its most significant restructuring since the field emerged as a distinct policy domain. In the span of roughly fourteen months, the Cybersecurity and Infrastructure Security Agency lost the majority of its workforce, the National Institute of Standards and Technology's AI Safety Institute was renamed and fundamentally reoriented, the foundational executive order on safe AI development was revoked on the first day of a new administration, and the Department of Defense—now styling itself the Department of War—released an AI acceleration strategy that explicitly deprioritizes compliance documentation in favor of fielding speed.

These changes do not represent isolated policy adjustments. Together they constitute a systemic shift in the federal government's posture toward AI risk: away from precautionary oversight, toward competitive deployment. The near-term security consequences are substantial. Enterprises deploying AI in regulated industries, critical infrastructure operators who previously relied on federal guidance, and government contractors navigating FedRAMP AI authorizations all face a landscape in which voluntary frameworks carry more practical weight than ever before—precisely as the agencies responsible for maintaining those frameworks are experiencing significant capacity reductions.

This whitepaper examines the specific capacity losses and policy redirections occurring across CISA, NIST, and the Department of Defense, traces their combined effect on the responsible scaling debate, and provides guidance aligned with the Cloud Security Alliance's frameworks for organizations seeking to maintain defensible AI security posture in the absence of strong federal coordination.

---

1. Introduction and Background

1.1 The Governance Architecture That Existed

For much of 2023 and 2024, US federal AI security governance rested on a coherent, if incomplete, architecture. Executive Order 14110, signed by President Biden in October 2023, established reporting requirements for foundation model developers, directed NIST to develop AI safety guidelines, and tasked sector regulators with issuing guidance for AI use in their domains. NIST responded by standing up the AI Safety Institute within its Information Technology Laboratory, a body designed to conduct safety evaluations, produce voluntary standards, and coordinate with international counterparts [1]. CISA, meanwhile, published an AI

Roadmap committing the agency to defending AI systems as part of its critical infrastructure mandate, and partnered with international partners to release joint guidance on securing AI in operational technology environments [2].

The NIST AI Risk Management Framework, released in January 2023 and updated through 2025, became one of the most broadly adopted voluntary AI governance standards in the United States [3]. Sector regulators including the CFPB, FDA, and SEC began referencing its principles in expectations for organizations deploying AI in regulated contexts. The overall architecture was voluntary and largely aligned across the relevant agencies: executive direction, NIST standard-setting, CISA operational guidance, and sector regulatory alignment.

That architecture began dissolving on January 20, 2025.

1.2 The Policy Rupture

President Trump's first-day Executive Order 14148 rescinded Executive Order 14110 in its entirety, eliminating the reporting requirements for foundation model developers and removing the legal basis for much of NIST's AI safety mandate [4]. A subsequent executive order, EO 14179, titled "Removing Barriers to American Leadership in Artificial Intelligence," established a new policy framework centered on deregulation and competitive advantage rather than safety precaution [5]. The order mandated creation of an AI Action Plan focused on maintaining US global dominance, with human flourishing and economic competitiveness named as the primary goals.

In December 2025, President Trump signed an additional executive order preempting state-level AI regulation, establishing a litigation task force to challenge state AI laws, and signaling that the federal government intended to occupy the field of AI governance—but in a form oriented toward enabling deployment, not constraining it [6]. The cumulative effect was to dismantle the Biden-era governance architecture while the replacement framework remained under construction.

The result is a structural vacuum: federal authority over AI security has been reasserted in principle while the specific institutions, mandates, and resources that gave it practical effect have been substantially reduced.

---

2. CISA: The Collapse of Operational AI Security Capacity

2.1 Workforce Reduction at Scale

The most concrete expression of the governance crisis is the condition of the Cybersecurity and Infrastructure Security Agency. CISA began fiscal year 2025 with approximately 3,400 employees and ended the year with roughly 2,400—a loss of approximately one thousand staff driven by a combination of DOGE-directed layoffs, voluntary buyouts, and early retirements [7]. By February 2026, with a partial federal government shutdown underway since February 14, CISA was operating with only 888 of its 2,341 active employees, placing the agency at approximately 38 percent of its functional staffing levels [8].

The proposed FY2026 budget compounds this trajectory. The Trump administration's budget request would reduce CISA's total funding from $2.87 billion in FY2024 to approximately $2.38 billion, with operational obligation levels falling by a further $420 million to approximately $1.96 billion [9]. The projected workforce for FY2026 stands at 2,324 full-time equivalent positions, compared with the pre-2025 level of approximately 3,292. The deepest cuts fall within risk management operations, where the division would decline from 179 positions to 58, and stakeholder engagement and requirements, from 200 positions to 53 [10].

2.2 Program Elimination and AI Implications

The human cost of these reductions is inseparable from the programmatic cost. The Election Security Program, which employed 14 staff and carried an annual budget of approximately $39.6 million, has been eliminated entirely [9]. The Cyber Defense Education and Training program faces a $45 million cut [9]. CISA's counter-ransomware initiative and its efforts to promote secure software development have both been degraded by the departure of the personnel who led them [11].

CISA's AI-specific work has not been formally discontinued, but its continuation is endangered by the staffing reductions documented above—particularly the reduction of the risk management operations division from 179 to 58 positions, the very function responsible for operationalizing sector-specific security guidance. The agency published meaningful AI security guidance in 2025, including a December 2025 joint publication with partners from Australia, Canada, Germany, the Netherlands, New Zealand, the United Kingdom, the NSA, and the FBI on principles for the secure integration of AI in operational technology environments [12]. This guidance addressed the accelerating adoption of machine learning, large language models, and AI agents across critical infrastructure, warning that integrating AI into systems that directly control physical processes introduces cyber, safety, and reliability risks that require specialized safeguards. The guidance's publication coincided with precisely the period in which the agency's capacity to operationalize and enforce its own recommendations was being dismantled.

May 2025 saw CISA release additional guidance for AI system operators on managing data security risks across the AI lifecycle, including securing the data supply chain and protecting data against unauthorized modification [13]. These publications represent genuine institutional knowledge. The concern is not that the documents were poorly conceived, but that the agency infrastructure needed to help critical infrastructure owners act on them is no longer intact.

2.3 The Cascade Effect on Private Sector and Partners

Federal cybersecurity capacity does not exist in isolation from the private sector ecosystem it supports. CISA's Joint Cyber Defense Collaborative, its threat information sharing programs, and its Cybersecurity Advisory products all depend on a staffed, funded agency maintaining operational tempo. When CISA's stakeholder engagement division goes from 200 staff to 53, the organizations that relied on that engagement for real-time threat context experience a commensurate reduction in federal support.

For AI security specifically, this matters because many critical infrastructure operators have reported limited internal AI security expertise as a challenge, meaning the most consequential AI deployment decisions in their environments depend on external authoritative guidance. They depended on CISA as the authoritative federal voice for sector-specific AI security guidance.

That voice has not gone silent, but it has been reduced to a fraction of its former volume at precisely the moment when the threat surface from AI deployment is expanding.

---

3. The NIST AI Safety Institute Rebranding and Mission Shift

3.1 From AISI to CAISI

In June 2025, Secretary of Commerce Howard Lutnick announced the transformation of the US AI Safety Institute into the Center for AI Standards and Innovation (CAISI) [14]. The rebranding was accompanied by a mission reorientation. Where AISI had been established with a focus on safety evaluations and the development of guidelines aligned with EO 14110's risk-management objectives, CAISI was chartered to "evaluate and enhance US innovation of these rapidly developing commercial AI systems while ensuring they remain secure to our national security standards" [14].

Secretary Lutnick's statement framing the change is worth examining carefully. He described the prior organization as having used "censorship and regulations" under "the guise of national security," positioning the rebrand as a liberation of innovators from arbitrary constraint. The practical implication is that CAISI's evaluation work will focus on "demonstrable risks, such as cybersecurity, biosecurity, and chemical weapons"—a narrower mandate than AISI's broader safety and trustworthiness orientation [14].

The institutional disruption accompanying this shift was significant. Elizabeth Kelly, AISI's first director, departed in early February 2025 [15]. DOGE-driven layoffs at NIST reportedly included many of the staffers who built AISI's technical evaluation capabilities [15]. The international network of AI Safety Institutes that AISI helped convene in San Francisco in November 2024, representing counterpart agencies from the UK, EU, Japan, and other partners, faces an uncertain future with its US founding institution now redirected.

3.2 Implications for Voluntary Standard-Setting

The NIST AI Risk Management Framework remains in place and continues to be updated. A March 2025 update broadened its coverage of generative AI and LLM vulnerabilities, added threat categories including poisoning attacks, evasion attacks, data extraction, and model manipulation, and emphasized model provenance and third-party model assessment [3]. Draft cybersecurity guidelines for the AI era were published by NIST in December 2025, with a full public draft anticipated in 2026 [16].

However, the institutional context in which these standards are developed and promulgated has changed. AISI's safety evaluation function—the practical work of testing frontier models for dangerous capabilities and providing independent assessments—is now subordinated to a "pro-innovation" mission. The combination of staffing reductions and mission reorientation means that the technical capacity for deep, independent AI safety evaluation within the US government has declined substantially. CAISI may continue to produce voluntary guidance and coordinate international standards work, but its role as an independent check on frontier AI development has been diminished by design.

4. The Department of Defense AI Acceleration Strategy

4.1 An AI-First Operating Model at Wartime Speed

The Department of Defense—which the Hegseth administration has rebranded as the Department of War, though the formal statutory name established by Congress remains unchanged—released its Artificial Intelligence Acceleration Strategy on January 9, 2026 [17]. The strategy, presented through two memoranda issued by Secretary Hegseth, represents the most operationally significant federal AI policy document of the current period. Its substance and its framing both deserve careful attention from security practitioners.

The strategy directs the Department toward an "AI-first operating model" measured by fielding speed rather than compliance documentation. It aims to make the United States "the world's undisputed AI-enabled fighting force" [17]. To this end, it commits to substantial investment in AI compute infrastructure from datacenters to the edge, mandates data access through "DoD Data Decrees" that give the Chief Digital and AI Officer authority to direct the release of any Department data to cleared users, and establishes a Chief Technology Officer-led innovation operating system designed to produce "next-generation technology, scalable products, and new ways of fighting" at "wartime speed" [18].

The FY2026 NDAA independently directs DoD to establish a cross-functional team led by the Chief Digital and AI Officer to create a standardized, Department-wide framework for assessing, governing, and approving AI model development, testing, and deployment, with the framework operational by June 1, 2026, and system assessments complete by January 1, 2028 [19].

4.2 The Redefinition of Responsible AI

Secretary Hegseth's public statement on the strategy included a definition of responsible AI that departs substantially from how that term has been understood in both commercial and prior federal policy contexts. According to Hegseth, "responsible AI at the War Department means objectively truthful AI capabilities employed securely and within the laws governing the activities of the department" [17]. He further specified that responsible AI means systems "free from ideological 'tuning'" with "no usage policy constraints for those systems beyond those imposed by statute" [17].

The effect of this definition, if not its stated intent, is to treat safety and alignment properties as political constraints rather than technical requirements—removing from the responsible AI concept the elements that have been its analytical core in civilian and academic contexts: robustness to adversarial manipulation, alignment with ethical principles, transparency and explainability, and proactive risk management. The definition focuses entirely on operational truthfulness and legal compliance.

The practical implications for enterprise AI procurement and security architecture are not trivial. Organizations building products and services for federal customers, particularly defense and intelligence community contracts, must navigate a client that has explicitly redefined its expectations for AI governance. The accountability structures and safety testing regimes that

were previously referenced in federal AI procurement are now at risk of being deprioritized in favor of deployment velocity.

4.3 FedRAMP's AI Authorization Acceleration

In parallel with the Department's acceleration strategy, the General Services Administration announced a FedRAMP 20x initiative in August 2025, prioritizing AI cloud service authorizations and accelerating the authorization process from months to weeks through automated validation [20]. Several major AI systems achieved FedRAMP authorization through this pathway in late 2025, including CrowdStrike Charlotte AI at FedRAMP High and Moveworks at FedRAMP Moderate [21, 22].

The acceleration of AI authorization processes reflects the broader deployment-first posture of the current federal technology policy environment. Faster authorization enables faster deployment; the question facing security practitioners is whether the accelerated timeline preserves the substantive security assessments that make FedRAMP authorizations meaningful, or whether it introduces new risks through insufficient evaluation depth.

---

5. The Responsible Scaling Vacuum

5.1 Industry Self-Regulation as the Primary Safety Mechanism

The retreat of federal oversight has elevated the importance of industry voluntary commitments to an unprecedented degree. Anthropic's Responsible Scaling Policy, version 3.0, became effective on February 24, 2026—one of the most detailed voluntary AI safety commitments currently operative in the United States [23]. Version 3.0 represents a significant evolution from its predecessors: it introduced Frontier Safety Roadmaps with detailed safety goals, and committed to publishing Risk Reports every three to six months explaining how capabilities, threat models, and active mitigations relate to each other [23].

The RSP framework introduced AI Safety Levels, with ASL-3 safeguards activated for relevant models in May 2025. However, a notable change in RSP v3.0 is the removal of the hard limit that barred Anthropic from training more capable models without safety measures already proven to work [24]. The updated policy separates two tracks: safety mitigations the company will pursue regardless of competitor behavior, and a broader capabilities-to-mitigations map recommending what the full industry should adopt. Critics noted that removing the hard commitment attenuates its deterrent value; Anthropic positioned the change as enabling more adaptive and transparent governance [24, 25].

The structural limitation of RSPs and similar voluntary commitments—including those from Google DeepMind and OpenAI—is precisely that they are voluntary. They bind only the signatory, they can be modified or withdrawn, and they are not subject to mandatory, binding external verification equivalent to formal regulatory audits—though Anthropic's RSP v3.0 does include provisions for independent third-party reviewers to evaluate unredacted Risk Reports under certain conditions [23]. When federal capacity for independent safety evaluation is simultaneously declining, as it is, the resulting architecture concentrates AI safety governance within the organizations that also have the strongest commercial incentives to deploy.

5.2 The Federal AI Action Plan Gap

EO 14179 directed completion of the AI Action Plan within 180 days of the order's signing in January 2025 [5]. As of early March 2026, no comprehensive federal AI Action Plan has been publicly released that establishes binding safety requirements or oversight mechanisms equivalent to those in the rescinded EO 14110. The December 2025 executive order focused on preempting state-level regulation without providing a substantive federal framework to replace what was removed. Organizations seeking to understand their federal compliance obligations regarding AI safety find themselves navigating a space defined primarily by what has been taken away rather than what has been put in its place.

The AI regulation landscape reflects this fragmentation. States have enacted significant AI legislation to fill the regulatory space vacated by the federal framework changes, with state-level AI legislation accelerating through 2025 and into 2026 [26]. The federal litigation task force established to challenge these state laws has created significant legal uncertainty for organizations operating across jurisdictions. The result is a patchwork in which enterprises face both the absence of coherent federal safety requirements and active federal opposition to the state-level frameworks that are attempting to provide them.

5.3 International Divergence

The United States' withdrawal from a safety-oriented AI governance posture is occurring against a backdrop of continued international standard-setting. The EU AI Act remains in force, with its high-risk AI provisions taking effect on August 2, 2026 [29]. The international network of AI Safety Institutes, though its US founding member has been reoriented, continues to coordinate on evaluation methodologies and risk frameworks. The UK AI Safety Institute remains active. This divergence creates compliance complexity for multinational organizations and technology companies operating across jurisdictions, which must satisfy EU safety requirements even as their home-country regulator has deprioritized comparable standards.

---

6. Risk Assessment: Enterprise and Sector Implications

6.1 Critical Infrastructure Operators

Critical infrastructure operators face the most acute governance gap. Sectors including energy, water, transportation, and healthcare have been expanding their AI deployments while the federal agency responsible for coordinating their security posture operates at 38 percent of optimal staffing. The OT-AI joint guidance published by CISA and international partners in December 2025 provides a useful technical baseline [12], but guidance without operational follow-through —sector-specific assessments, threat intelligence sharing, incident coordination—delivers substantially reduced value.

Organizations deploying AI in systems that directly control physical processes carry risk that cannot be fully addressed through voluntary framework adoption alone. The absence of federal operational support raises the baseline cost and complexity of maintaining AI security in these environments, while simultaneously increasing the likelihood that industry-specific bad actors may attempt to exploit the oversight gap.

6.2 Federal Contractors and the FedRAMP Landscape

Federal contractors deploying AI systems into government environments must navigate the tension between the administration's deployment-acceleration posture and the substantive security requirements that remain embedded in FedRAMP and CMMC frameworks. Faster authorization timelines create business opportunity while potentially compressing the evaluation depth on which customers rely. Organizations seeking to build defensible security postures should treat FedRAMP authorization as a floor rather than a ceiling, maintaining AI-specific security controls beyond what the authorization process requires.

The NDAA's mandate for a DoD AI governance framework by June 2026 may provide some structure, but its focus is on the Department's internal AI deployments rather than on the broader contractor ecosystem. Contractors should anticipate that the resulting framework will reflect the DoD's operational definitions of responsible AI rather than the broader safety and trustworthiness standards that governed prior procurement.

6.3 Commercial AI Deployers in Regulated Industries

For commercial organizations deploying AI in regulated industries—financial services, healthcare, pharmaceuticals—the governance vacuum creates paradoxical risk. Sector regulators including the SEC, CFPB, and FDA have been referencing NIST AI RMF principles in their expectations for AI deployment [3]. These references remain operative even as the federal safety architecture around NIST has been weakened. Organizations that reduce their AI governance investments in response to the federal framework changes may find themselves exposed when sector regulators enforce existing supervisory expectations.

The liability framework is also shifting. As Biden-era safety mandates are revoked and state requirements face federal preemption challenges, fewer prescriptive regulations tell organizations exactly what AI security measures to implement. This shifts the analytical burden toward negligence standards: organizations that cannot demonstrate reasonable care in AI deployment decisions may face litigation exposure for failing to meet an evolving reasonable-care standard, particularly as the lowered regulatory floor makes the absence of proactive governance more visible [26].

---

7. Conclusions and Recommendations

7.1 Strategic Assessment

The United States is in a period of AI governance transition that is likely to persist for years, with substantive resolution unlikely before a new governance architecture is established to replace what was dismantled. The Biden-era architecture has been dismantled; a replacement that provides comparable clarity on safety requirements has not yet materialized. The responsible scaling commitments of frontier AI developers represent the most substantive operative safety governance in the current environment, but they are voluntary, concentrated, and subject to competitive pressure.

Federal capacity for AI security—through CISA's operational role and NIST's standard-setting role—has been reduced substantially from pre-2025 levels. The combination of workforce reduction at CISA, mission reorientation at CAISI, and the DoD's operational redefinition of responsible AI creates a governance environment in which enterprises must do substantially more of their own AI security work than the pre-2025 architecture anticipated.

This is not an argument for or against any particular policy orientation. It is an assessment of the operational reality facing security practitioners and risk managers in 2026. Organizations that plan their AI security governance assuming a level of federal support that no longer exists will carry unpriced risk.

7.2 Recommendations for Enterprise Organizations

The recommendations below are organized around the specific gaps that the current governance transition has created. Each addresses a category of AI security responsibility that organizations can no longer reasonably delegate to federal oversight capacity, and each is grounded in frameworks and resources that remain available regardless of the policy environment.

Internalize federal guidance that may not be operationally supported. CISA's December 2025 OT-AI joint guidance and its May 2025 AI data security guidance remain technically sound and represent the best available federal baseline for AI security in operational contexts [12, 13]. Organizations should use these documents as frameworks even as CISA's capacity to operationalize them has declined.

Treat NIST AI RMF adoption as a primary risk management instrument. The March 2025 NIST AI RMF update, with its expanded coverage of generative AI threats and emphasis on model provenance and third-party assessment, provides actionable risk management structure that is not contingent on federal enforcement [3]. Adoption aligned with this framework provides both substantive security value and a defensible record for sector regulatory purposes.

Monitor responsible scaling commitments from AI providers. Anthropic's RSP v3.0, including its Frontier Safety Roadmaps and quarterly Risk Reports, represents a new level of voluntary transparency that enterprise procurement teams should incorporate into AI vendor assessment [23]. OpenAI's and Google DeepMind's equivalent commitments should be evaluated comparably. When assessing these commitments, procurement teams should evaluate not only current content but also commitment durability—voluntary policies can be modified, and the removal of RSP v3.0's hard capability limits illustrates how the terms of these commitments can shift. The quality, specificity, and stability of these commitments now constitute a meaningful differentiator in enterprise AI procurement risk analysis.

Prepare for EU AI Act obligations regardless of domestic policy posture. Organizations with any EU market exposure face binding high-risk AI requirements effective August 2026 [29]. These obligations are not affected by US federal deregulation. Organizations that have deferred EU AI Act compliance work based on the US regulatory environment should accelerate their preparation.

Increase AI security investment specifically in areas that CISA has vacated. Threat intelligence sharing, sector-specific AI security assessments, and incident response planning for AI system failures are areas where CISA's reduced capacity creates the most acute gaps for

critical infrastructure operators. Organizations in these sectors should consider consortium arrangements with peers and increased direct engagement with ISACs to partially substitute for the federal coordination role.

7.3 Recommendations for Policy and Governance Stakeholders

The CSA urges federal policymakers to recognize that the governance transition underway carries structural risks distinct from the policy debates it is occurring within. Whatever the appropriate balance between AI safety and AI competitiveness, the current state—in which the institutions responsible for AI security guidance are operating well below functional capacity while AI deployment in critical infrastructure accelerates—represents a period of elevated systemic risk that merits urgent attention.

Specific policy actions that would reduce this risk without prejudging the broader governance debate include: restoring CISA's operational capacity to a level commensurate with the AI security threats facing critical infrastructure; providing CAISI with sufficient technical staffing to maintain meaningful independent evaluation capability; and publishing the AI Action Plan mandated by EO 14179 with sufficiently specific AI security standards that regulated industries have a coherent federal baseline to reference.

CSA Resource Alignment

The issues analyzed in this whitepaper connect directly to several CSA frameworks and working group outputs that provide operational guidance for organizations navigating the current governance environment.

The **MAESTRO framework** (Multi-Agent Environment Security Threat and Risk Orchestration) addresses the security architecture of agentic AI systems of the type now being deployed at accelerating pace in both federal and commercial environments. As CISA's AI security capacity declines, the threat modeling structures in MAESTRO become more, not less, relevant as a first-line risk management tool for organizations deploying multi-agent AI systems without federal operational support.

The **AI Organizational Responsibilities series**, including the volume on Governance, Risk Management, Compliance and Cultural Aspects, provides RACI models and implementation strategies for organizations establishing internal AI governance programs [27]. In a governance vacuum, these internal structures carry heightened importance. The series addresses shadow AI, model drift monitoring, incident response planning, and board-level AI oversight—precisely the governance domains that organizations can no longer delegate to federal oversight capacity.

The **AI Risk Management: Thinking Beyond Regulatory Boundaries** whitepaper from the CSA AI Governance and Compliance Working Group addresses the limitations of compliance-only approaches to AI risk management [28]. Its argument—that organizations should pursue proactive risk management that exceeds minimum compliance requirements—is especially timely in an environment where the minimum compliance floor is being reduced. The document's emphasis on trustworthiness, transparency, and lifecycle risk management provides a durable governance orientation that is not contingent on any particular regulatory architecture.

The **Cloud Controls Matrix (CCM) v4** provides specific control mappings for AI systems that can be applied to enterprise procurement, vendor assessment, and internal AI deployment governance. Organizations using CCM as their primary cloud security control framework should apply its principles to AI system procurement and deployment decisions, particularly as FedRAMP AI authorization processes are accelerated.

The **CSA AI Controls Matrix** and its companion auditing guidelines provide a structured basis for evaluating AI systems against security and governance requirements. As voluntary responsible scaling commitments from frontier AI developers become the primary operative safety governance, enterprise customers using these tools to evaluate provider commitments will be better positioned to make informed procurement decisions.

References

[1] National Institute of Standards and Technology, "Center for AI Standards and Innovation (CAISI)," NIST, 2025, https://www.nist.gov/aisi [Note: URL now serves CAISI content; original AISI page has been replaced by the successor organization.]

[2] Cybersecurity and Infrastructure Security Agency, "Roadmap for AI," CISA, 2023, https://www.cisa.gov/resources-tools/resources/roadmap-ai

[3] ISPartners, "NIST AI RMF 2025 Updates: What You Need to Know," ISPartners Blog, March 2025, https://www.ispartnersllc.com/blog/nist-ai-rmf-2025-updates-what-you-need-to-know-about-the-latest-framework-changes/

[4] AO Shearman, "President Trump Rescinds Biden's AI Executive Order," AO Shearman Insights, January 2025, https://www.aoshearman.com/en/insights/ao-shearman-on-tech/president-trump-rescinds-bidens-ai-executive-order-and-announces-ai-investment

[5] Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence," The White House, January 23, 2025, https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/

[6] Mayer Brown, "President Trump Issues Executive Order on 'Ensuring a National Policy Framework for Artificial Intelligence'," Mayer Brown Insights, December 2025, https://www.mayerbrown.com/en/insights/publications/2025/12/president-trump-issues-executive-order-on-ensuring-a-national-policy-framework-for-artificial-intelligence

[7] Axios, "One-Third of U.S. Cyber Agency CISA Has Left Since Trump Took Office," Axios, June 3, 2025, https://www.axios.com/2025/06/03/cisa-staff-layoffs-resignations-trump-cuts

[8] TechCrunch, "US Cybersecurity Agency CISA Reportedly in Dire Shape Amid Trump Cuts and Layoffs," TechCrunch, February 25, 2026, https://techcrunch.com/2026/02/25/us-cybersecurity-agency-cisa-reportedly-in-dire-shape-amid-trump-cuts-and-layoffs/

[9] Federal News Network, "DHS Budget Request Would Cut CISA Staff by 1,000 Positions," Federal News Network, May 2025, https://federalnewsnetwork.com/cybersecurity/2025/05/dhs-budget-request-would-cut-cisa-staff-by-1000-positions/

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

[11] Valley Techlogic, "Government Backed Cybersecurity Agency CISA Down to Just 38% of Its Optimal Staffing Levels," Valley Techlogic, February 27, 2026, https://www.valleytechlogic.com/2026/02/27/government-backed-cybersecurity-agency-cisa-down-to-just-38-of-its-optimal-staffing-levels-after-funding-cuts-what-it-means-for-your-business/

[12] Cybersecurity and Infrastructure Security Agency, "Principles for the Secure Integration of Artificial Intelligence in Operational Technology," CISA, December 2025, https://www.cisa.gov/resources-tools/resources/principles-secure-integration-artificial-intelligence-operational-technology

[13] Inside Privacy, "CISA Releases AI Data Security Guidance," Inside Privacy, May 2025, https://www.insideprivacy.com/cybersecurity-2/cisa-releases-ai-data-security-guidance/

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

[15] Fortune, "AI Safety Advocates Slam Trump Administration's Reported Targeting of Standards Agency," Fortune, February 20, 2025, https://fortune.com/2025/02/20/trump-doge-layoffs-nist-aisi-ai-safety-concerns/

[16] National Institute of Standards and Technology, "Draft NIST Guidelines Rethink Cybersecurity for the AI Era," NIST News, December 2025, https://www.nist.gov/news-events/news/2025/12/draft-nist-guidelines-rethink-cybersecurity-ai-era

[17] Department of Defense, "War Department Launches AI Acceleration Strategy to Secure American Military AI Dominance," Press Release, January 12, 2026, https://www.war.gov/News/Releases/Release/Article/4376420/war-department-launches-ai-acceleration-strategy-to-secure-american-military-ai/

[18] Defense One, "The Pentagon Says It's Getting Its AI Providers on 'the Same Baseline'," Defense One, February 2026, https://www.defenseone.com/technology/2026/02/pentagon-says-its-getting-its-ai-providers-same-baseline/411506/

[19] WilmerHale, "What the NDAA Means for AI and Cybersecurity," WilmerHale Insights, December 19, 2025, https://www.wilmerhale.com/en/insights/client-alerts/20251219-what-the-ndaa-means-for-ai-and-cybersecurity

[20] General Services Administration, "GSA and FedRAMP Announce Major Initiative: Prioritizing 20x Authorizations for AI Cloud Solutions," GSA News, August 25, 2025, https://www.gsa.gov/about-us/newsroom/news-releases/gsa-fedramp-prioritize-20x-authorizations-for-ai-08252025

[21] CrowdStrike, "CrowdStrike Charlotte AI Achieves FedRAMP High Authorization," CrowdStrike Press Release, November 2025, https://ir.crowdstrike.com/news-releases/news-release-details/crowdstrike-charlotte-ai-achieves-fedramp-high-authorization

[22] ServiceNow Newsroom, "Moveworks from ServiceNow Achieves FedRAMP Moderate Authorization," ServiceNow, February 2026, https://newsroom.servicenow.com/press-releases/details/2026/Moveworks-from-ServiceNow-achieves-FedRAMP-moderate-authorization-to-provide-secure-conversational-AI-to-public-sector/default.aspx

[23] Anthropic, "Responsible Scaling Policy Version 3.0," Anthropic, February 24, 2026, https://anthropic.com/responsible-scaling-policy/rsp-v3-0

[24] Winbuzzer, "Anthropic Drops Hard Safety Limits From its AI Scaling Policy," Winbuzzer, February 25, 2026, https://winbuzzer.com/2026/02/25/anthropic-drops-hard-safety-limit-responsible-scaling-policy-xcxwbn/

[25] Open Data Science, "Anthropic Updates Responsible Scaling Policy to Strengthen AI Risk Governance," Open Data Science, February 2026, https://opendatascience.com/anthropic-updates-responsible-scaling-policy-to-strengthen-ai-risk-governance/

[26] Corporate Compliance Insights, "2026 Operational Guide to Cybersecurity, AI Governance and Emerging Risks," Corporate Compliance Insights, 2026, https://www.corporatecomplianceinsights.com/2026-operational-guide-cybersecurity-ai-governance-emerging-risks/

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

[28] Cloud Security Alliance, "AI Risk Management: Thinking Beyond Regulatory Boundaries," CSA AI Governance and Compliance Working Group, 2024

[29] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 on Artificial Intelligence (EU AI Act)," Official Journal of the European Union, July 2024, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689