March 31, 2026

ISSUE BRIEF | AI and Emerging Technology

BUILDING AI READINESS IN THE U.S. GOVERNMENT

Cole Salvador, Jack Crovitz, & Yusuf Mahmood

TOPLINE POINTS

- **The federal government is not fully prepared for the AI revolution**. Despite bold steps by the Trump Administration, further actions are needed by Congress and executive agencies. This issue brief proposes a concrete AI readiness agenda across three areas: talent, adoption, and strategic foresight.
- **We must accelerate the deployment of AI talent into government.** Pay caps, slow hiring, and divestment rules make government uncompetitive with industry. Congress should fund the U.S. Tech Force, create a part-time Tech Force Reserve, and establish a "Highly Qualified AI Experts" hiring authority.
- **We must accelerate AI adoption in government.** Dysfunctional procurement processes and classified compute shortages delay or kill AI adoption initiatives. Congress should enable colorless AI acquisition, expand Other Transaction Authority government-wide, direct the publication of agent security standards, codify the CAIO network, and direct the rapid expansion of secure AI compute capacity for sensitive use.
- **We need hubs of strategic AI foresight in government.** The Center for AI Standards and Innovation and the Bureau of Emerging Threats must be authorized to execute the Administration's AI priorities and to guard their current missions against political turnover. Congress should increase their funding and give them interagency levers.

Introduction

For years, the U.S. government has been dangerously behind on AI. Many trends shaping AI today, such as exponentially rising energy demand, were foreseeable years in advance. A small team of informed analysts could have predicted them. But the federal government lacked the staff, mandate, and technical capacity to see them coming and to prepare appropriate responses.

President Trump has moved decisively to reassert American AI leadership in the private sector. The administration has unleashed innovation by repealing burdensome Biden-era regulations and DEI requirements. The nation has begun clearing energy bottlenecks through executive action, Federal Energy Regulatory Commission (FERC) rule changes, and the National Energy Dominance Council. AI adoption at the Pentagon is accelerating. The administration has reformed federal procurement to block ideologically biased AI products.

Yet the White House's *AI Action Plan* recognizes that the federal government has a unique and ongoing role to play in winning the AI race, a role that the private sector cannot fill. Its essential functions include:

(1) **Horizon-scanning and proactive deregulation.** The federal government must engage in horizon-scanning, i.e., understanding where AI is heading and preparing to facilitate the technology's progress. Horizon-scanning can also aid with proactive deregulation: eliminating cumbersome and unnecessary regulations *before* our innovators face critical shortages. For example, the energy bottleneck impacting the AI boom could have been anticipated years earlier—and deregulation initiated sooner—if a small office had been tracking compute scaling laws and extrapolating obvious trends.
(2) **National security uses.** Only the federal government can and should comprehensively assess AI systems for capabilities relevant to national security, evaluate foreign AI, and coordinate classified threat assessments across agencies.
(3) **Procurement and adoption.** Only the federal government can purchase AI systems for government use-cases (i.e., procurement) and adopt AI throughout the federal government. The federal government should capture the same efficiency gains from AI that the private sector enjoys.

The federal government is structurally unable to meet these needs. Federal pay scales top out around $250,000—a fraction of the compensation that private industry offers, even to median AI talent. Hiring processes are slow and bureaucratic. Procurement rules and technology regulations make it difficult to adopt cutting-edge AI tools rapidly. The centers of technical capacity that the U.S. government has designated, such as the Center for AI Standards and Innovation (CAISI) and the State Department's Bureau of Emerging Threats (ET), are not authorized by statute or appropriately funded. The Biden Administration also tried to co-opt similar institutions by externally imposing harmful DEI initiatives; future administrations may try again. As a result, the government risks falling so far behind the private sector that it can no longer play its essential role in winning the AI race.

To stay ahead, the government should pursue three crucial priorities:

(1) **Hiring talented AI experts** with technical expertise to know: (i) what AI systems agencies need to procure; (ii) how to integrate them into government systems; and (iii) how the latest AI capabilities should inform broader AI strategy.

(2) **Accelerating Federal AI adoption** to make federal agencies more agile, lethal, and ready for the AI revolution that is already reshaping private enterprises.

(3) **Building strategic foresight.** With AI talent and adoption in place, offices can be empowered to fill the government's role in the AI race. These offices should receive authorization and levers to understand key trends in AI before they become strategically critical.

Part I: Hiring Talented AI Experts

Advanced AI tools can help the federal government better serve the American people, from defending our borders and stopping cartels from trafficking fentanyl, to detecting fraud and eliminating government waste—but the federal workforce needs AI talent to realize that potential.

We discuss three barriers that stand in the way of the federal government's ability to hire top AI talent rapidly: (1) federal hiring processes are burdensome; (2) public sector wages are not competitive for AI roles; and (3) AI talent is dissuaded by low-prestige jobs and DEI-infused processes.

We then describe several solutions to these problems, including: (1) authorizing the U.S. Tech Force; (2) creating a "Tech Force Reserve" that can make use of technical talent without interrupting private-sector careers; (3) leveraging Office of Personnel Management (OPM) authorities to boost salaries for technical staff; (4) expanding "excepted service" hiring; and (5) promoting flexible hiring mechanisms in the federal government.

Problem: The Federal Government Cannot Attract High-Quality AI Talent

Problem 1: Federal hiring processes are burdensome. Federal technical staffing is broken because of its uniquely dysfunctional hiring processes. Hiring regulations for traditional "competitive service" positions require rigid ranking and quantitative selection, interview quotas, and inflexible recruitment processes. These rules, originally designed to combat favoritism and relationship-based hiring, now merely stall the process and enable discrimination through DEI programs (Executive Office of the President, 2021). Furthermore, career bureaucrats have extensive legal protections that make firing difficult and risky. These protections, which were enhanced under the Biden Administration to make at-will firing of federal employees impossible, discourage agencies from hiring except when overwhelmingly necessary (White House, 2025a).

To combat some of these inefficiencies, Congress has authorized the Office of Personnel Management (OPM) to extend "Direct-Hire Authority" (DHA) to specific agencies or for specific roles with a "severe shortage of candidates" or "critical hiring need" (eCFR, 5 §337). This authority allows for hiring without regard to most of the complex administrative requirements established by U.S. law. OPM has explicitly authorized the use of DHA for technical roles relating to AI through 2028 (Ahuja, 2023). According to public job listings, agencies hiring for technical AI positions often make use of DHA (for example: USA JOBS, "Senior Cyber Offense Specialist;" "AI Standards Architect").

Though DHA makes it easier to select candidates during a hiring process, attracting suitable candidates is a greater challenge.

Problem 2: Public sector wages are not competitive for AI roles. Today, AI experts are perhaps the most in-demand workers in history. Silicon Valley's AI labs offer excellent salaries to secure the best among them. Meta reportedly offers pay packages as high as $1.5 billion to new members of its Superintelligence initiative (Shibu, 2025). The AI boom's thirst for talent is not limited to AI model developers, either. At semiconductor design firm Nvidia, for instance, over 80% of staff have become millionaires (Gairola, 2025). In contrast, the 2025 General Schedule (GS) — the pay scale determining federal employee pay for non-executives — disallows salaries over $195,200 (OPM, 2025a). The Executive Schedule, which applies to a small number of positions (most of which must be confirmed by the Senate), is capped at $250,600 (OPM, 2025b). These maximum salaries are

not competitive with industry wages, where even median software engineers earn over $500,000 (Levels.fyi, n.d.).

Figure 1
Median Pay at AI Companies vs. Maximum Federal Salary

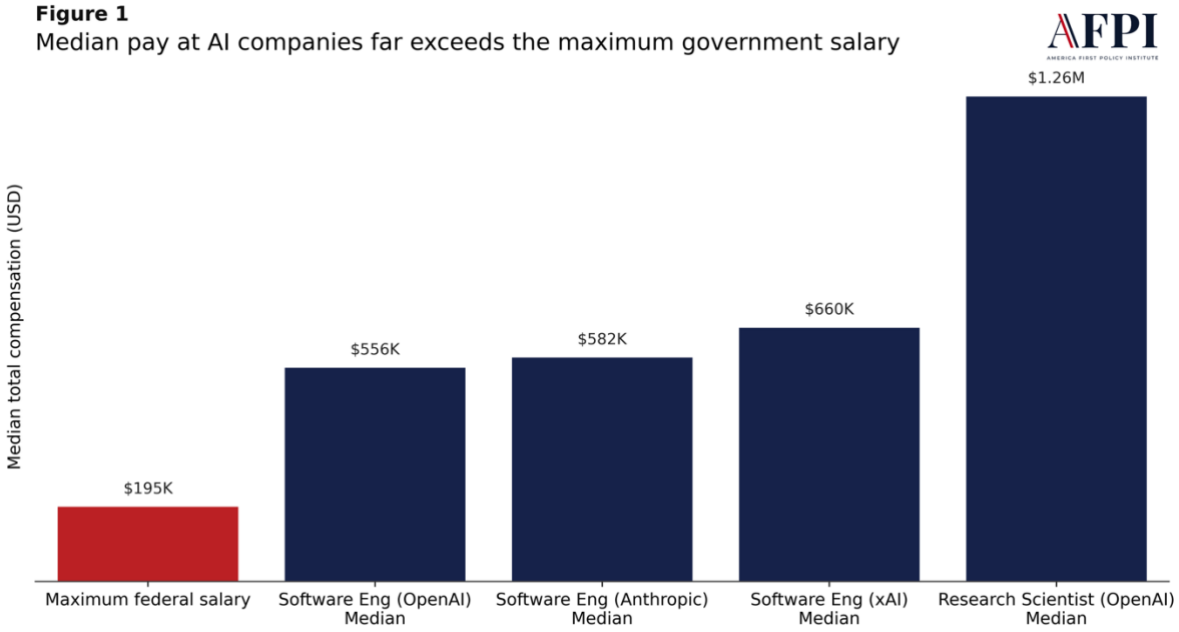

Note. The maximum federal salary for non-executive positions, which includes virtually all technical roles, is not competitive with private sector salaries in the AI industry. Data from Office of Personnel Management, *Salary Table 2025-DCB,* 2025, and Levels.fyi, n.d.

Problem 3: Technical experts do not want to enter the complacent, DEI-infused, and frustrating federal bureaucracy. In the AI industry, many top researchers choose their employer because of their ideological allegiance, mission focus, and a "startup" environment (Burleigh, 2025; Shibu, 2025). Working in the federal government has long been seen as a non-prestigious, undesirable career (Hattery, 1955). For many prospective employees, it does not feel like an agile organization or a smart career move. In 2018, for example, tech specialists in government over the age of 60 outnumbered those under 30 by four times (Mazur, n.d.). Workers also view government as a non-meritocratic system: only 47% believe that "differences in performance are recognized in a meaningful way" in their work (Shriver, 2025).

Further, AI experts with private sector experience are often hesitant to join the government because it would mean forgoing private sector income and divesting most of their equity in private tech companies (eCFR, 5 §2635.402). But these things can change. As the Trump Administration dismantles the bloated administrative state, critical offices in technology and national security may become attractive opportunities for technologists to serve their country.

Solutions to the AI Talent Problem

Solution 1: Authorize the U.S. Tech Force. In December 2025, the U.S. Office of Personnel Management (OPM), in collaboration with various federal agencies, announced the U.S. Tech

Force. The program consists of cohorts of 1,000 fellows serving one- to two-year terms in technical roles at federal agencies (Kupor, 2025).

Hiring for Tech Force is centralized by OPM. This allows the program to take advantage of economies of scale in hiring, which it has done. For example, the announcement of Tech Force was accompanied by a media blitz, high-profile support and outreach, and the release of a dedicated website (National Design Studio, n.d.). All these are possible because of the program's scale. These efforts have largely succeeded: Tech Force received over 35,000 expressions of interest in its first month (Bracken, 2026). Tech Force fellows will be onboarded directly into agencies, "with teams of about 30-40 individuals at most large agencies" (Kupor, 2025). They will be managed by Tech Force leaders drawn from the private sector, not by career bureaucrats. Agencies are fully responsible for Fellow salaries. Congress could appropriate $50 million to the program, allowing OPM to pay for one-third of each Fellow's salary to incentivize technical hiring.

Solution 2: Create a "Tech Force Reserve," designed to make use of technical talent without interrupting private sector careers. Reserve Fellows could be designated Special Government Employees (SGEs) under existing authority, permitting them to work for 130 days or less, a year in the federal government, and without requiring them to leave the private sector (DOJ, 2006). The Secretary of War has already directed the Pentagon to use the Tech Force program to expedite hiring under its AI Acceleration Strategy (Department of War, 2026).

Congress could allow the Department of War to pilot such a program as it adopts AI beyond administrative tasks. Congress could authorize $100 million to the DOW program for compensation and security clearances. In the short-term, Secretary Hegseth could implement a smaller pilot program for tech-focused reservists under SGE authority.

Solution 3: The Office of Personnel Management (OPM) can boost technical salaries. Agencies can offer yearly awards up to $10,000 to employees, and with approval from OPM, these awards can extend to $25,000 (eCFR, 5 §451). Additionally, agencies can offer annual staff bonuses not to exceed 25% of salary, or 50% with OPM approval. OPM sets a firm limit on total pay, including awards and bonuses, at the highest level of the executive schedule, so General Schedule employees can receive a maximum of $250,600 per year, including any awards and bonuses (OPM, n.d.-a). OPM could also extend "critical position pay" to certain AI roles, allowing them to exceed that maximum (OPM, n.d.-b).

Solution 4: Agencies can expand the use of flexible hiring authorities. Flexible hiring authorities allow agencies to bring on a surge of new talent for a limited period, avoiding typical red tape. Options for expansion include the following:

- **Agencies can expand the use of the Intergovernmental Personnel Act (IPA) program.** IPA allows agencies to temporarily hire staff from universities, federally funded research and development centers (FFRDCs), nonprofits, and other organizations engaged in activities of mutual concern with the federal government (OPM, n.d.-b). The cost burden of IPA staff can be shared between the sponsoring organization and receiving agency (OPM, n.d.-b).
- **OPM can expand the use of Schedule A and Schedule B "excepted service" hiring.** These authorities are generally exempt from federal hiring procedures and allow for open recruiting of specific candidates (eCFR, 5 §302). Using these authorities, OPM could exempt specific agencies — or AI-related positions more generally — from "competitive service" regulations.

Solution 5: Congress could expand the existing Department of War Public-Private Talent Exchange (PPTE) program. Since the 2017 National Defense Authorization Act (NDAA), the PPTE program has allowed DOW to hire term-limited personnel from the private sector. PPTE was designed as a complement to the Intergovernmental Personnel Act (IPA), which is limited to academia and FFRDCs. The intense concentration of AI talent in the private sector warrants an expansion of the PPTE program to AI-focused roles across the federal government, not just in the Pentagon.

Solution 6: Congress could explore legislation that establishes novel hiring exemptions for critical AI-related positions. For example, Congress could introduce a "Highly Qualified AI Experts" program based on the Pentagon's existing "Highly Qualified Experts" (HQE) hiring authority, which allows the Secretary of War to hire a limited number of experts without following federal hiring regulations and to authorize additional payments for these HQEs outside the regular pay range for comparable positions (H.R. 1588, 2003). This "Highly Qualified AI Experts" program could authorize agency heads or Chief AI Officers to hire up to 20 AI experts with the same procedural and financial flexibility that the Department of War enjoys with HQEs.

Part II: Accelerating AI Adoption in the Federal Government

President Trump has made government AI adoption a priority of his administration. As the White House's *AI Action Plan* declared, "With AI tools in use, the Federal government can serve the public with far greater efficiency and effectiveness" (White House, 2025b). AI adoption will allow federal agencies to maximize their efficiency, responsiveness, and lethality. Conversely, if the American government fails the AI adoption challenge, it will fall behind private enterprises and foreign adversaries, undermining national security and public trust in government services.

Unfortunately, the federal government historically lags the private sector in adopting new technologies. Federal agencies report that even when AI tools are ready for off-the-shelf use, rapid acquisition and adoption are not feasible due to bureaucratic hurdles, lack of leadership, and a shortage of secure AI compute. The Trump Administration has already taken concrete steps to address these issues. Further action, including legislation, can empower agencies to adopt state-of-the-art commercial AI tools to improve mission outcomes and serve the American people.

Streamlining Procurement of Commercial AI Tools

Problem: Many agencies report that federal procurement systems and requirements create lengthy delays in AI adoption initiatives or prevent them entirely. Traditional federal procurement processes are complex and highly regulated, and they often represent stumbling blocks for agencies that attempt to integrate cutting-edge AI tools into their workflows (Cooper, 2025). These delays can be "exacerbated when the provider is unfamiliar with federal procurement requirements"—like most commercial AI developers (GAO, 2025a). The need to follow complex federal procurement processes also often discourages businesses from attempting to conduct business with the federal government at all (Exec. Order 14275, 2025).

Solution 1: Congress could establish a "colorless" acquisition process for AI procurement in the Pentagon. Effective military adoption of AI is among the most urgent national security imperatives of the 21st century. As the White House's *AI Action Plan* noted, "The United States must aggressively adopt AI within its Armed Forces if it is to maintain its global military preeminence" (White House, 2025b). Secretary of War Hegseth has pushed the Pentagon to embrace commercial AI tools and streamline AI procurement (Hegseth, 2025a; Hegseth, 2025b), but some inefficiencies are imposed by statute.

One example is the "color of money" system, which requires Pentagon expenses to be categorized as (1) research, (2) procurement, or (3) maintenance. This system conflicts with modern commercial software development cycles, in which all three activities occur simultaneously. Software systems are continuously improving even while deployed and in active use. This mismatch can make it nearly impossible for the Department of War to buy or license commercial software. AI and software acquisitions are routinely delayed for months—or even cancelled—because they seem to require three different "colors of money" (McQuade et al., 2019; Waterman, 2025). Such interruptions are particularly destructive in AI, which advances so quickly that even a few months' delay could render an acquisition obsolete.

Since 2019, the "Single Appropriation Pilot for Software and Digital Technology Budget Activity," also known as BA-08, has allowed the Pentagon to use small amounts of research-categorized funds for software procurement and maintenance. The pilot has been an important tool for the Department of War to invest in cutting-edge technology, demonstrating that "colorless" software acquisition is operationally valuable (Serbu, 2023; Auchey, 2024). Purchasing AI tools using "colorless"

funds would enable the Pentagon to rapidly adopt state-of-the-art AI tools to maximize efficiency and lethality.

Solution 2: Congress could expand Other Transaction Authority (OTA) for AI procurement across the federal government. OTA is a procurement authority that allows the Department of War to acquire technology with more flexibility than regular procurement regulations permit. Secretary Hegseth has made OTA the "default" tool for military acquisition of AI tools, encouraging procurement professionals to use it to make rapid, commercial-style transactions (Hegseth, 2025a). Congress has gradually extended OTA to a handful of civilian agencies — for example, the Advancing American AI Act of 2022 expanded OTA to DHS (H.R. 7776, 2023, §7221-7228) — but many agencies still lack the authority. Further expanding agencies' access to OTA would accelerate AI adoption across the federal government.

Solution 3: The Government Accountability Office (GAO) could survey agencies concerning regulations and statutes that currently serve as barriers to the rapid procurement and adoption of AI tools. This survey would build on previous investigations (for example: GAO, 2025a), and could inform Congressional priorities for further reform

Developing a Security Framework for AI Agent Deployment

AI agents — autonomous systems that can browse the web, execute code, manage files, and take real-world actions with minimal human oversight — offer enormous potential to make federal agencies faster, leaner, and more lethal (GAO, 2025b).

Problem: Some federal agencies are delaying or abandoning AI agent adoption initiatives because they lack clear, authoritative frameworks for deploying these systems securely. Agencies understand that AI agents introduce security considerations that are qualitatively different from those posed by chatbots or traditional software (Chambers, 2026). NIST's Center for AI Standards and Innovation (CAISI) reports that AI agent security concerns in U.S. government agencies "hinder adoption today" (Press, 2026). For example, the Chief AI Officer of the Department of Labor notes that AI agents' unique capabilities require heightened safety protocols compared to other AI tools (Heckman, 2026). Some of these concerns are legitimate. Attackers can, for example, hijack AI agents by embedding malicious instructions in unassuming emails, documents, or web pages (Datta et al., 2025). Research has shown that attackers can hijack agents to execute unauthorized code, exfiltrate sensitive files, and send targeted phishing messages using a victim's own contact list (CAISI, 2025). These vulnerabilities are already being exploited in production systems across the private sector (for one particularly dramatic example, see Reddy & Gujral, 2025).

Solution: The National Institute of Standards and Technology (NIST) can encourage AI agent adoption by developing practical agent security guidance. A growing body of research shows that automated system-level controls can mitigate federal agencies' security concerns with AI agents. Clear guidance from NIST could empower agencies to implement these security protocols and thereby accelerate their deployment of AI agents. Critically, NIST's guidance should not be one-size-fits-all. Low-stakes administrative workflows require different protections than intelligence or defense applications. By offering tiered frameworks calibrated to different deployment contexts, NIST can give agencies the confidence to move forward with AI agent adoption at a pace appropriate to their mission rather than defaulting to inaction.

Empowering and Training a Network of AI Adoption Leaders

Problem 1: Federal agencies need embedded individuals and bodies that take responsibility for accelerating AI adoption, breaking down barriers to AI acquisition, and sharing best practices across the federal government. Federal AI adoption has long remained limited due to siloed agency AI initiatives and the lack of a cross-agency network of AI adoption leaders (Elbaum & Hippold, 2025). Agencies need "key officials to lead agency AI adoption and promote the sharing of best practices," but existing law generally does not account for such officials (Vought, 2025). Adoption moves too slowly if agencies waste time duplicating each other's work.

Solution: Congress could pass legislation to codify the Trump Administration's institutional framework for AI adoption leadership. Together, OMB's memo M-25-21 and the White House's *AI Action Plan* described a network of AI leaders across the federal government (Vought, 2025; White House, 2025b). This network includes:

(1) **Chief AI Officers (CAIOs)**. Each agency is mandated to appoint a CAIO, an official with "significant expertise in AI" who is responsible for promoting AI adoption within the agency.

(2) **The Chief AI Officer Council (CAIOC)**. This is an interagency council, led by OMB and consisting of the CAIOs of all the federal agencies as well as representatives from OSTP, ODNI, and other expert offices. Its mission is to enable interagency coordination on AI adoption, including the sharing of data, technical resources, and best practices.

(3) **AI Governance Boards**. Each agency is mandated to regularly convene an internal AI Governance Board to support the CAIO's initiatives, coordinate AI policy within the agency, and explore the governance requirements of their agency's use of AI.

Most of these positions and bodies are not yet authorized by statute. To ensure consistency across agencies and prevent political turnover from disrupting this system of interagency collaboration, Congress could confirm the existence and duties of all these roles and bodies in legislation.

Problem 2: Much of the federal procurement workforce does not understand the state of AI technology, commercial AI acquisition practices, or the need for rapid AI adoption in the federal government. Agencies report that a lack of expertise in AI technology and commercial software acquisition is among the most pressing obstacles to AI adoption (GAO, 2025a). Congress has already tried to train procurement professionals for AI adoption with the AI Training Act of 2022, which directed OMB to prepare annual training programs for procurement professionals to better understand the technology (S. 2551, 2022). This initiative is not enough. In 2024, the OMB program had under 15,000 registrations, even though there are 186,000 procurement staff in the Pentagon alone (GSA, 2024; Gates et al., 2022). The curriculum also does not include training in commercial software acquisition practices, which will be essential expertise for efficiently acquiring AI tools.

Solution: Congress could update the AI Training Act. First, Congress could improve participation rates by tying completion of OMB's AI training program to existing Federal Acquisition Certification requirements. Simultaneously, Congress could add commercial software acquisition practices as a required topic of the annual training program. The Pentagon's Defense Innovation Unit (DIU) would be a useful partner in this curriculum as it is already preparing training materials on that topic (DIU, 2025). These amendments would transform the AI Training Act into an effective vehicle for familiarizing the federal acquisition workforce with both AI technology and commercial software contracting practices.

Expanding AI Compute Infrastructure for Sensitive AI Workflows

Problem: Many federal agencies state that access to trusted, high-security AI computing power is a major bottleneck in their ability to use AI for tasks involving classified or otherwise sensitive data. The Department of War and the intelligence community (IC) face this challenge most acutely (GAO, 2025). As AI adoption accelerates and agencies increasingly fine-tune AI models for national security use-cases, abundant access to high-security AI compute will become a crucial pillar of the U.S. government's AI strategy.

The Trump Administration recognizes the federal government's urgent and increasing need for high-security AI compute. As the *AI Action Plan* explains, "It is likely that AI will be used with some of the U.S. government's most sensitive data. The data centers where these models are deployed must be resistant to attacks by the most determined and capable nation-state actors" (White House, 2025b).

Unfortunately, the infrastructure for such high security "classified compute" barely exists today. Total AI computing infrastructure has surged over the past few years, but the vast majority of that capacity sits in commercial data centers with no security accreditation. The supply of high-performance GPU compute inside classified environments is vastly smaller than the AI compute available for commercial uses. As one analyst notes, "the U.S. national security enterprise lacks anything comparable to private companies' capacity in AI infrastructure" (Bajraktari, 2024; Bajraktari, 2025).

Figure 2

Companies Have Rapidly Increased Their Share of AI Supercomputer Ownership

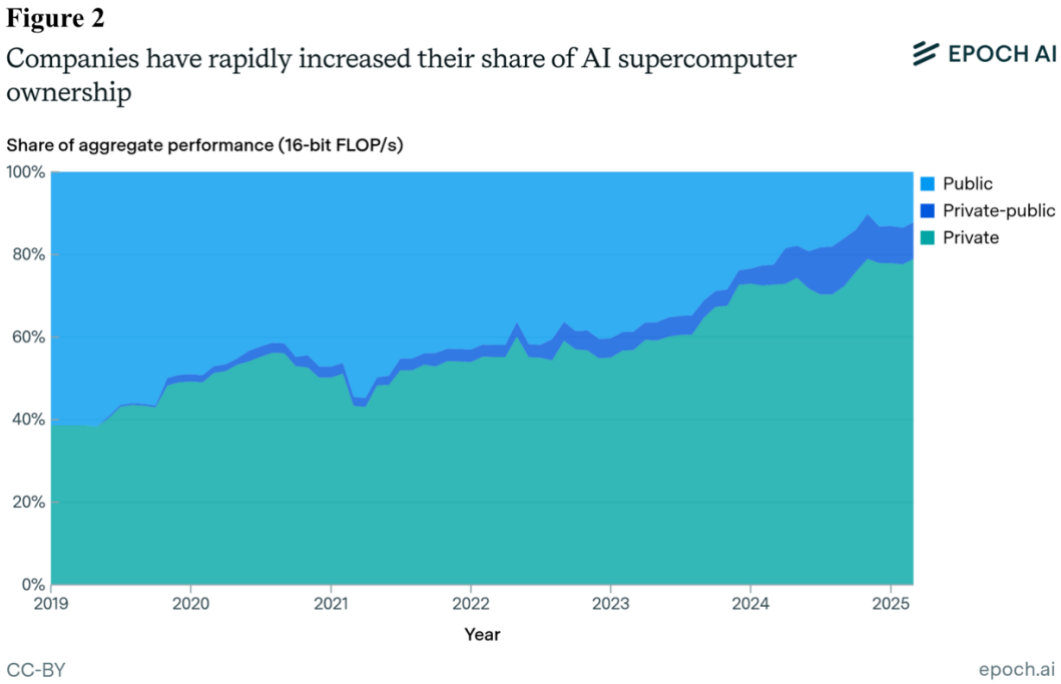

Note. Public and private AI supercomputer ownership. The "Public" category includes academic or other nongovernmental initiatives, and only a small share of "Public" or "Private-public" computing power is in classified environments. Figure from: Pilz, K. et al., *Data on GPU Clusters*, 2026.

One particularly stark example: In 2025, CENTCOM had a classified AI compute shortage, so it worked with the Pentagon's Chief Digital and AI Office (CDAO) to acquire 28 NVIDIA

H100 GPUs. The chief data officer of CENTCOM boasted that this acquisition "will give us a really huge, significant amount of compute capability that no one else — at least that I'm tracking — has in the Defense Department for classified networks" (Obis, 2025). In contrast, that same year, commercial hyperscalers routinely built AI data centers with tens or hundreds of thousands of NVIDIA H100 GPUs (Epoch AI, 2026).

As the Pentagon's demand for specialized compute for AI inference and fine-tuning grows, the classified compute shortage will become more severe. We must ensure that American warfighters do not fall behind private actors and adversaries in AI adoption due to a lack of classified compute.

Solution: The U.S. government should prepare for expanded demand for classified compute by proactively expanding the supply. The Pentagon is already drafting the "technical standards for new secure datacenters" as a result of the *AI Action Plan* (Hegseth, 2026), and the Department of Energy is constructing AI data centers as part of the Genesis Mission (DOE, 2025). Congress could direct DOE and the Department of War to lead a cross-agency effort to construct or retrofit AI data centers certified to handle classified information.

We suggest certain parameters for this initiative:

(1) **Public-Private Partnership.** The data center construction or retrofitting should be completed in partnership with established private hyperscalers, perhaps with ownership or revenue-sharing interest in the data center's operation. This public-private partnership would allow the government to leverage corporate expertise while defraying costs to the taxpayer. DOE's planned "Solstice" AI data center (which it is constructing with NVIDIA and Oracle and is predicted to come online in 2026) is a valuable precedent for such partnerships (DOE, 2025).

(2) **Security Level.** The data centers should operate at a level of information security that would allow them to prevent infiltration or exfiltration by advanced nation-state adversaries (Nevo, 2025). They should be accredited to process classified information up to the TS/SCI level.

(3) **Compute Scale.** The inter-agency initiative should aim to construct or retrofit data centers with a scale comparable to the most recent commercial facilities, including at least 100,000 H100 GPUs or equivalent computing capacity. This capacity matches the DOE's planned "Solstice" AI data center (DOE, 2025).

(4) **Timeline.** The initiative should aim to retrofit or construct secure data centers on an ambitious timeline — perhaps as brief as 18 months. This timeline is not unheard of in commercial data center construction (Somala et al., 2025).

The resulting secure data centers would serve as shared national resources available to DOW, the IC, and other agencies with classified AI workloads. This initiative would be merely the first step in supplying the U.S. government with high-security AI compute. Once operational, the facilities could serve as proof of concept and reference architecture for further public or private efforts to expand high-security AI compute capacity, including facilities for fine-tuning national security AI models.

Part III: Building Hubs of Strategic Foresight on AI

American AI dominance requires in-house talent that can understand and anticipate the technology. This is especially urgent given AI's rapid pace of progress: academic papers become obsolete shortly after publication, and every month brings a new best model.

In this section, we discuss two possible hubs of strategic foresight for AI: the Center for AI Standards and Innovation (CAISI) in the Department of Commerce and the Bureau of Emerging Threats (ET) in the Department of State.

CAISI has two barriers to success: (1) lack of funding and staff, and (2) lack of a focused mission. We argue that it should receive substantially more funding. We then discuss an America First vision for CAISI's mission, in which it functions as technical strike team, bridge between industry and government, frontier analysis unit, and technical standards organization.

ET has two barriers to success: (1) lack of congressional authorization and (2) lack of interagency influence. We recommend that Congress authorize the Bureau to deliver national security insights and give it concrete levers of interagency influence.

Strategic Foresight on AI is Valuable

Strategic foresight a few years ago could have allowed the U.S. government to prepare for the challenges that jeopardize American AI leadership today. Consider the rise of deep learning and the energy bottlenecks we now face. To attentive technologists, these trends were predictable: deep learning was a natural consequence of trends in semiconductor performance (LeCun et al., 2015); advanced large language models (LLMs) were identified as a consequence of "compute scaling laws" formulated as early as 2019 (Kaplan et al., 2020); and the energy and infrastructure demands of AI products were predictable at least as early as 2022, if not earlier (Sevilla et al., 2022). Straight lines on graphs could have told us the future.

The U.S. government failed to anticipate technological trends and their strategic importance because it lacked the staff, mandate, and capacity to understand AI. Even a small office dedicated to analyzing AI progress (such as Epoch AI, which analyzes machine learning trends) could have foreseen the coming energy bottlenecks and AI's importance years in advance (Sevilla et al., 2022; Owen, 2025). Building this capacity in government is essential.

The Center for AI Standards and Innovation

The Center for AI Standards and Innovation (CAISI), housed within the National Institute of Standards and Technology (NIST) at the Department of Commerce, shows promise as a hub of federal AI expertise. It has conducted leading model evaluations and analysis of international AI competition, such as its report on the performance and censorship of leading Chinese and U.S. AI models (CAISI, 2025). But CAISI has often been hampered by a lack of a clear mission or appropriate funding. Through targeted action, CAISI could become the country's touchpoint for accelerating the private sector and ensuring we stay ahead of our adversaries.

Today, CAISI has talented staff, most of whom are computer scientists, engineers, and machine learning experts with experience at frontier AI companies, research universities, and startups. It also has memoranda of understanding and non-disclosure agreements with top AI developers,

including OpenAI and xAI, for model evaluation and research collaboration; and with multiple foreign countries, including the United Kingdom, for standards development and talent exchange (White House, 2025c).

Problem 1: CAISI Lacks Funding and Staff

CAISI is chronically underfunded. Since its founding in 2024, it has received only $30 million in total, not annual funding (U.S. Senate Committee on Appropriations, 2024). It does not have any regular annual appropriations. Analogous AI institutes in other countries have received larger appropriations, often with smaller mandates: Singapore's at $39 million (Singapore AISI, 2024); Canada's at $36 million (Canadian AISI, 2024); and the United Kingdom's over $330 million (UK Department for Science, Innovation & Technology, 2026).

Figure 3

U.S.'s CAISI Funding Compared to Peer Institutes

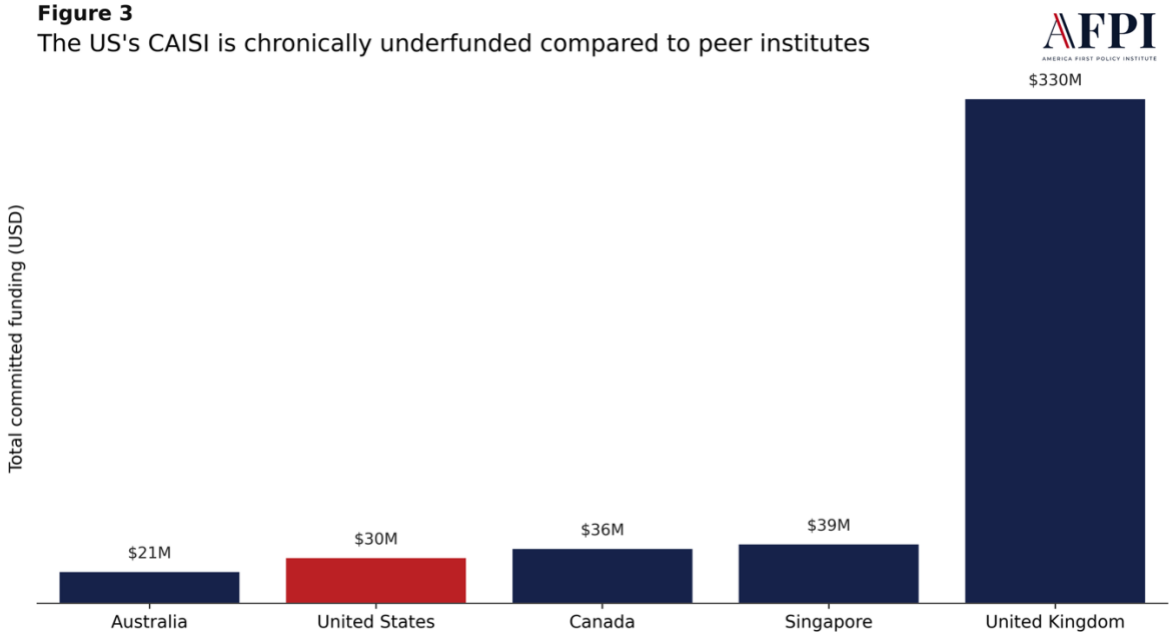

Note. Committed funding (**total, not annual**) for the U.S. Center for AI Standards and Innovation (CAISI) lags far behind analogous institutes in other countries and is over ten times smaller than in the United Kingdom. Data from: (Singapore AISI, 2024; Canadian AISI, 2024; UK Department for Science, Innovation & Technology, 2026; Charlton, 2025; U.S. Senate Committee on Appropriations, 2024).

CAISI's lack of funding is especially damaging because the organization depends on costly computing resources for evaluations, research, and other missions. As a result, its headcount-to-budget ratio is smaller than most agencies.

CAISI is also chronically understaffed. With only 20-30 full-time employees, it has more missions than people (HKS, 2026). Its staff includes technical researchers and engineers and national security experts with a technology focus, but its small size limits its capacity to fulfill its many missions and, increasingly, to field inbound requests from other agencies like the IC.

Solution 1: Congress should fund CAISI with a $50-100 million annual budget. CAISI is clearly among the most promising hubs of AI expertise in government. America should not devote fewer resources to AI foresight than peers like Australia, Canada, and the UK. To grow CAISI and let it execute its critical missions, Congress should fund it with an annual appropriation of $50-100 million.

Solution 2: Let CAISI hire quickly. The Office of Personnel Management (OPM) should ensure CAISI has preferential access to Tech Force Fellows and the hiring authorities as previously described.

Problem 2: CAISI Does Not Have a Focused Mission

CAISI has a huge volume of missions without Congressional authorization. To date, the Center has publicly received taskings from two sources. First, Secretary of Commerce Howard Lutnick announced in June 2025 (Lutnick, 2025) that CAISI will:

- Develop guidelines and best practices on the security of AI systems, in collaboration with NIST;
- "Establish voluntary agreements" with the frontier AI industry and lead evaluations of demonstrable national security risks;
- Assess the state of international AI competition, the capabilities of U.S. and adversary systems, and the adoption of foreign systems;
- Evaluate models for security vulnerabilities, such as backdoors;
- Coordinate interagency actions on AI evaluations and assessments, including with DOW, DOE, DHS, OSTP, and the IC; and
- "Represent U.S. interests internationally" to promote American AI standards and defend against unnecessary foreign regulation.

Second, the White House's *AI Action Plan* tasked CAISI with several activities (White House, 2025b). The *AI Action Plan* directed CAISI to lead or support actions across evaluations, standards development, national security, and research and development. Table 1 summarizes these tasks.

Table 1
Summary of Recommended Policy Actions that the White House's AI Action Plan *Tasked to CAISI*

AI Action Plan: Recommended CAISI Actions	
Category	**Tasking**
Evaluations	Publish research and evaluations of Chinese AI models for CCP biases (p. 4)
	Evaluate frontier AI systems for national security risks, especially in chem-bio (p. 22)
	Evaluate U.S. and foreign AI systems, and U.S. critical infrastructure, for foreign influence (p. 22)
	"Build, maintain, and update" national security-related AI evaluations (p. 22)
Standards development	Issue guidance for Federal agencies to conduct bespoke AI evaluations (p. 10)
	Convene biannual meetings on AI evaluation best practices (p. 10)
	Collaborate with national security agencies to create technical standards for high-security AI data centers (p. 16)
	Assist in publishing an AI assurance standard for the IC (p. 18)
National security	Collaborate with agencies to collect intelligence on foreign frontier AI projects (p. 6)
	Collaborate with AI developers to protect against malicious foreign and insider threats (p. 12)
	Assist in establishing an AI Information Sharing and Analysis Center (p. 18)
	Partner with industry to include AI in incident response team planning (p. 19)
	Modify CISA cybersecurity response playbooks to account for AI systems (p. 19)
Research and development	Collaborate with DARPA and NSF to develop AI interpretability, control, and adversarial robustness (p. 9)
	Coordinate an AI hackathon initiative with academia and various agencies (p. 10)

CAISI's missions could be threatened by future externally imposed DEI initiatives. Under the Biden Administration, the Center sometimes had to navigate externally imposed Diversity, Equity, and Inclusion (DEI) initiatives that threatened to derail its work. In 2023, for example, CAISI's parent agency created an "AI Risk Management Framework" that instructed AI developers how to confront "harmful biases" against "demographic groups" or "individuals with disabilities" even in the absence of intent (NIST, 2023). A 2024 Presidential National Security Memorandum also tried to task CAISI with voluntary safety testing for "discrimination and bias, … and the safety of individuals and groups" (Biden, 2024).

Solution: An America First Vision for CAISI

The Center for AI Standards and Innovation (CAISI) needs a clear mission and the resources to execute it. In the table below (Table 2), we summarize the high-level roles we recommend for CAISI. We also list the goal of each role and a few example activities in that category. None of these lists are meant to be comprehensive. Congress should mandate CAISI to fulfill these roles while leaving enough flexibility to fulfill emerging roles without new statutory language. It should also explicitly bar CAISI from DEI-related activities that future administrations could impose on it.

Table 2
Overview of Core Priorities AFPI Recommends for the Center for AI Standards and Innovation

AFPI Recommended Priorities for CAISI		
Role	**Goal**	**Example Activities (Not Comprehensive)**
Technical Strike Team	Provide hands-on technical AI capacity to government agencies	Accelerate AI agent adoption by solving security challenges
		Evaluate frontier AI systems for national security risks, especially in chem-bio and cyber
		Provide technical support to America's AI Exports Program
Bridge Between Industry and Government	Help industry win the AI race through government support	Act as a "front door" for industry to request government help, such as on deregulation or counter-distillation
		Coordinate interagency processes to implement cross-cutting government AI priorities
		Collaborate with the IC to share threat intelligence with industry and develop threat models
Frontier Analysis Unit	Monitor the state of strategic competition between the US and its adversaries to deliver strategic AI foresight	Issue a quarterly report to senior national security leaders on the state of international AI competition
		Receive voluntary, confidential reports from industry on the future of AI development and competition
		Monitor AI adoption in government and global markets
Technical Standards Organization	Develop and promote AI usage, security, and measurement guidance	Issue guidance to government agencies on bespoke AI evaluations and adoption goals
		Combat harmful foreign regulations in international fora and standards organizations
		Conduct research on AI metrology and security standards

(a) CAISI as a technical strike team. Though small, CAISI is the government's densest concentration of AI talent. It should deploy that capacity to benefit America's AI goals wherever it is needed most. CAISI's unique mix of technical staff with national security expertise, for example, could be used to help AI developers (whose staff are mostly foreign nationals) develop insider threat programs. This role imagines CAISI in part as in-house technical consultants for agencies and private-sector developers implementing important AI programs.

(b) CAISI as a bridge between industry and government. AI does not fit neatly into the missions of most existing agencies. Consider, for example, how the government might counter the illegal "distillation attacks" Chinese AI companies perpetrate against U.S. developers to steal model capabilities (OpenAI, 2026). The Intelligence Community (IC) has the authority to monitor the attackers and disrupt their operations but lacks clear motivation or frequent touchpoints with industry to do so.

In cases like these, CAISI could act as a "front door" for industry to express its priorities and develop cohesive strategies that help American companies win the AI race. It could activate these activities and help coordinate them through interagency processes.

(c) CAISI as a frontier analysis unit. As we have argued, America's most pressing AI policy challenges—like energy and competition from China—were foreseeable. We should strengthen CAISI's horizon-scanning functions to preempt tomorrow's challenges. CAISI must be designed to: (1) foresee the most important future trends in AI impacting U.S. national security and technological leadership and (2) communicate these trends to key policymakers. In these roles, it would contribute to ensuring "that the appropriate agencies within the national security enterprise possess sufficient technical capacity to understand frontier AI," as the White House National Policy Framework for AI recommends (White House, 2026).

To make CAISI a frontier analysis unit, Congress could mandate it to do the following:

- Receive voluntary, confidential reports on the future of AI competition and development from U.S. AI companies, through existing and future memorandums of understanding (MOU);
- Help inform IC collection priorities for foreign AI development in adversary states;
- Analyze the above to develop insights on AI development and competition; and
- Issue a quarterly report, with classified annex, to key stakeholders on the state of AI development and competition, including APNSA, DNI, NEDC, Cabinet Secretaries, House and Senate Intelligence and China committee leadership, etc.

(d) CAISI as a technical standards organization. Secretary Lutnick correctly identified CAISI as an ideal standards development and promotion organization within government. It should develop AI standards on government adoption, data, evaluations, security, and measurement science with input from the National Institute of Standards and Technology (NIST). CAISI and NIST should also promote these standards throughout government and the world.

CAISI can leverage its measurements of federal AI adoption to ensure agencies adhere to important AI guidance like the Office of Management and Budget (OMB)'s 2025 Memorandum on Accelerating Federal Use of AI (Vought, 2025). Further, it can use its standards and measurement work to combat harmful foreign regulations that affect American companies, such as the EU AI Act. This work would play out in international fora like the United Nations and groups like the International Organization for Standardization.

Congress should authorize and fund CAISI specifically for national security and objective metrology, lest future administrations co-opt it for DEI initiatives. Though CAISI has not worked on DEI initiatives, leaving it unauthorized risks future administrations co-opting it to advance harmful woke ideologies. CAISI should promote human flourishing through objective analysis of AI, not act as a pawn of far-left ideologues who seek to weaponize the technology against Americans.

The Bureau of Emerging Threats

The Bureau of Emerging Threats (ET) is an agency established in 2025 by Secretary Marco Rubio's reorganization of the Department of State (Rubio, 2025). ET is "charged with anticipating and responding" to "U.S. adversaries' weaponization of advanced technology, including artificial intelligence" (Kingston, 2026). This mission is critical as state-sponsored cyber threat actors in Iran, North Korea, China, and Russia begin to use AI in "all stages of their operations" (GTIG, 2026). The Bureau reportedly has five divisions: Cybersecurity, Critical Infrastructure Security, Disruptive Technology, Space Security, and Threat Assessment.

The precursor to ET, the Bureau of Cyberspace Security and Emerging Technologies (CSET), was established in 2021 under President Trump. Then, the State Department requested that CSET receive a $20.8 million budget and a staff of 80 employees (Bugos, 2021). Although little is known about ET's current operations, we estimate its current staffing and funding levels are smaller, given the recent reorganization. It nevertheless has significant expertise in military uplift and international analysis relating to AI.

Problem: ET Lacks Authorization

ET is not authorized by statute. The reorganization of the State Department implemented under Secretary Rubio in 2025 has not yet been codified in law. This means that Congress does not directly fund the Bureau. It also means that ET lacks protection from interference by future administrations, which might deemphasize its mission or co-opt the agency for harmful Diversity, Equity, and Inclusion initiatives. Future administrations could also dismantle the Bureau entirely, as President Biden did to the 1776 Commission (Conklin, 2021).

The Bureau has significant talent in international relations, including career diplomats and national security experts (Bureau of Emerging Threats, n.d.). It also has a strong relationship with the Department of State's Intelligence Community component, the Bureau of Intelligence and Research (INR), which has extensive knowledge of our adversaries. But uncertainty created by a lack of authorization likely has a chilling effect on hiring and funding for the Bureau. This challenges its ability to deliver insights to policymakers.

Solution 1: Congress should authorize ET. Authorization of ET and the entire State Department reorganization would codify the "America First State Department" built by Secretary Rubio that reduces administrative bloat and focuses its mission (Rubio, 2025). Congress should mandate ET to deliver national security insights on emerging technology to key policymakers. Congress should also prevent ET from being co-opted by future administrations to engage in DEI initiatives.

Solution 2: The State Department should bolster ET's interagency influence. As a unique hub of government AI foresight, ET needs levers to communicate its insights throughout government. The State Department, or Congress, where appropriate, should give the bureau

these levers — like those we recommend for the Center for AI Standards and Innovation (CAISI). See "CAISI as a frontier analysis unit" above for some of the levers we suggest.

Policy Recommendations

The following recommendations, if implemented, could ensure that the federal government plays its proper role ensuring U.S. AI dominance. We summarize our recommendations around the three categories discussed in this issue brief: (1) hiring talented AI experts; (2) accelerating AI adoption; and (3) building hubs of strategic foresight on AI.

Hiring Talented AI Experts

OPM and agencies should aggressively utilize legal pathways to expedite the hiring of AI technical experts and expand their pay. These include Schedule A and Schedule B "excepted service" hiring authorities; the Intergovernmental Personnel Act (IPA) program; and awards under CFR §451, which require OPM approval.

Congress should pass legislation to empower agencies to attract and rapidly onboard AI technical talent. Options include:

- Authorizing and funding the current "U.S. Tech Force" program with $50 million;
- Establishing a "U.S. Tech Force Reserve" program to designate a limited number of private-sector AI technical experts as Special Government Employees;
- Expanding the existing Department of War Public-Private Talent Exchange (PPTE) program to civilian agencies; and
- Establishing a "Highly Qualified AI Experts" hiring authority that would authorize agencies to hire a limited number of AI experts with significant procedural and financial flexibility.

Accelerating AI Adoption

Congress should pass legislation to streamline AI procurement processes. Options include:

- Establishing a "colorless" acquisition process for AI procurement in the Department of War;
- Expanding Other Transaction Authority (OTA) across the government to allow agencies to acquire commercial AI products with the speed and agility of private enterprises; and
- Tasking the Comptroller General with issuing a report on existing statutes that prevent or delay AI adoption in the federal government.

Congress should direct NIST to publish clear AI agent security frameworks to give agencies the confidence they need to deploy the technology.

Congress should pass legislation to empower an interagency network of AI adoption leaders to drive their agencies to rapidly adopt commercial AI tools. Options include:

- Confirming the authority and duties of the AI adoption network established by OMB memo M-25-21, including agency Chief AI Officers (CAIOs), the Chief AI Officer Council, & agency AI Governance Boards; and
- Amending the AI Training Act of 2022 to tie completion of OMB's annual AI training program to existing Federal Acquisition Certification requirements and to cover commercial software acquisition.

Congress, the Department of War, and the intelligence community should rapidly expand classified AI compute infrastructure, both for inference and fine-tuning.

Building Hubs of Strategic Foresight on AI

Congress should authorize hubs of strategic AI foresight. The Department of Commerce's Center for AI Standards and Innovation (CAISI) should serve a focused mission and be allocated $50-100 million annually to serve these roles. Congress should also authorize the Department of State's Bureau of Emerging Threats (ET) to deliver national security insights and to provide it with levers of interagency influence. Both should be barred from DEI initiatives to prevent meddling by future administrations.

Conclusion

America's decisive advantage in the race for global AI leadership is clearly the private sector. But winning the larger strategic competition will require the federal government to be ready to facilitate AI development, prepare for AI-related national security risks, and take full advantage of the power of the technology to make America more agile, resilient, and lethal. This technology is far too commercially, strategically, and militarily important for the U.S. government to be left in the dark.

If our recommendations are implemented by Congress and the executive branch, the federal government would be more prepared to navigate the AI revolution. Our policies would enable the federal government to better understand AI trends, adopt AI for the benefit of our taxpayers and national security, and ensure continued U.S. AI dominance.

Works Cited

Ahuja, K. (2023, December 29). *OPM Memorandum on Government-wide Hiring Authorities for Advancing Federal Government Use of Artificial Intelligence (AI)*. https://www.opm.gov/chcoc/transmittals/2023/Government-wide%20Hiring%20Authorities%20for%20Advancing%20Federal%20Government%20Use%20of%20Artificial%20Intelligence%20%28AI%29%2012-29-2023.pdf

Auchey, J. (2024, May 28). *Pilot is paramount.* U.S. Army. https://www.army.mil/article/276687/pilot_is_paramount

Awards. 5 C.F.R. §451 (2026). https://www.ecfr.gov/current/title-5/chapter-I/subchapter-B/part-451

Bajraktari, Y. (2024, December 12). *Protecting the U.S. AI Compute Advantage*. Special Competitive Studies Project (SCSP). https://scsp222.substack.com/p/protecting-the-us-ai-compute-advantage

Bajraktari, Y. (2025, March 14). *FR Doc. 2025-02305: Input from the Special Competitive Studies Project on the Request for Information on the Development of an Artificial Intelligence (AI) Action Plan*. https://files.nitrd.gov/90-fr-9088/SCSP-AI-RFI-2025.pdf

Biden, J. (2024, October). *National Security Memorandum on Advancing the United States' Leadership in Artificial Intelligence*. https://www.presidency.ucsb.edu/documents/national-security-memorandum-advancing-the-united-states-leadership-artificial

Bracken, M. (2026, January 14). *OSTP's Kratsios touts 'incredible' interest in Tech Force, defends Trump science and tech cuts*. FedScoop. https://fedscoop.com/trump-ai-action-plan-michael-kratsios-tech-force

Bugos, S. (2021, April). *State Reviews Plans for New Tech Bureau*. Arms Control Today. https://www.armscontrol.org/act/2021-04/news/state-reviews-plans-new-tech-bureau

Bureau of Emerging Threats. (n.d.). *Leadership – Bureau of Emerging Threats*. Retrieved March 16, 2026, from https://www.state.gov/leadership-bureau-of-emerging-threats

Burleigh, E. (2025, August 4). *Anthropic CEO Dario Amodei says his employees are refusing Zuckerberg's $100 million payout—and he's not even matching salaries to keep them*. Fortune. https://fortune.com/2025/08/04/billionaire-anthropic-ceo-dario-amodei-ai-staffers-poaching-meta-mark-zuckerberg-100k-six-figure-salaries-openai-sam-altman

Chambers, A. (2026). *Request for Information Regarding Security Considerations for Artificial Intelligence Agents*. National Institute for Standards & Technology. https://www.govinfo.gov/content/pkg/FR-2026-01-08/pdf/2026-00206.pdf

Conklin, A. (2021). *Biden to rescind Trump's '1776 Commission'*. Fox News. https://www.foxnews.com/politics/biden-rescind-trump-1776-commission

Cooper, M. (2025, September 29). *AI Won't Outrun Bad Procurement*. RAND Research & Commentary. https://www.rand.org/pubs/commentary/2025/09/ai-wont-outrun-bad-procurement.html

Dahl, K. R. (2007, July). *New Security for New Threats: The Case for Reforming the Interagency Process*. The Brookings Institution. https://www.brookings.edu/wp-content/uploads/2016/06/dahl20070731.pdf

Datta, S., Nahin, S., Chabra, A., & Mohapatra, P. (2025). *Agentic AI Security: Threats, Security, Evaluations, and Open Challenges*. *arXiv*. https://arxiv.org/pdf/2510.23883v1

Defense Innovation Unit. (2025, July). *Immersive Commercial Acquisition Program (ICAP)*. https://www.diu.mil/work-with-us/immersive-commercial-acquisition-program-icap

Disqualifying Financial Interests. 5 C.F.R. §2635.402 (2026). https://www.ecfr.gov/current/title-5/chapter-XVI/subchapter-B/part-2635/subpart-D/section-2635.402.

Elbaum, S. & Hippold, M. (2025, August 5). *AI in the Federal Government: A Fragmented Reality*. Council on Foreign Relations. https://www.cfr.org/articles/ai-federal-government-fragmented-reality

Elysee. (n.d.). *AI Action Summit*. Retrieved March 16, 2026, from https://www.elysee.fr/en/sommet-pour-l-action-sur-l-ia

Epoch AI. (n.d.). *Frontier Data Centers*. Retrieved March 16, 2026, from https://epoch.ai/data/data-centers

Examining System. 5 C.F.R. §337 (2026). https://www.ecfr.gov/current/title-5/chapter-I/subchapter-B/part-337#5:1.0.1.2.48.2.16.1

Exec. Order No. 14275 (2025, April 15). *Restoring common sense to federal procurement*. https://www.whitehouse.gov/presidential-actions/2025/04/restoring-common-sense-to-federal-procurement/

Fist, T. & Datta, A. (2024, October 23). *How to Build the Future of AI in the United States*. Institute for Progress. https://ifp.org/future-of-ai-compute/#the-ai-data-centers-of-the-future

Gairola, A. (2025, August 5). *Nvidia Is Producing 'Unprecedented Wealth' For Its Employees, Nearly 80% Are Already Millionaires: Report*. Yahoo News. https://finance.yahoo.com/news/nvidia-producing-unprecedented-wealth-employees-003124691.html

Gates, S., Roth, E., and Kempf, J. (2022, November 28). *Department of Defense Acquisition Workforce Analyses*. RAND. https://www.rand.org/pubs/research_reports/RRA758-2.html

Google Threat Intelligence Group (GTIG). (2026, February 12). *GTIG AI Threat Tracker: Distillation, Experimentation, and (Continued) Integration of AI for Adversarial Use*. Google Cloud Blog. https://cloud.google.com/blog/topics/threat-intelligence/distillation-experimentation-integration-ai-adversarial-use

Government of Canada. (2024, November 12). *Canada launches Canadian Artificial Intelligence Safety Institute*. https://www.canada.ca/en/innovation-science-economic-development/news/2024/11/canada-launches-canadian-artificial-intelligence-safety-institute.html

Haley, V. & Ezell, C. (2025, May 29). *OPM Memorandum on the Merit Hiring Plan.* Office of Personnel Management. https://www.opm.gov/chcoc/transmittals/2025/Merit%20Hiring%20Plan%205-29-2025%20FINAL.pdf

Harvard Kennedy School Institute of Politics. (2026, January 5). *Director's Internship: Center for AI Standards and Innovation (CAISI)*. Retrieved [March 10, 2026] from https://iop.harvard.edu/internships-and-careers/directors-internship/host-organizations/center-ai-standards-and-innovation

Hattery, L. (1955). The Prestige of Federal Employment. *Public Administration Review, 15*(3). https://www.jstor.org/stable/973015

Heckman, J. (2026, March 23). *AI boosts efficiency for agencies, but trust and safety lead the way*. Federal News Network. https://federalnewsnetwork.com/artificial-intelligence/2026/03/ai-boosts-efficiency-for-agencies-but-trust-and-safety-lead-the-way/

Hegseth, P. (2025, March 6). *Memorandum on Directing Modern Software Acquisition to Maximize Lethality*. https://media.defense.gov/2025/Mar/07/2003662943/-1/-1/1/DIRECTING-MODERN-SOFTWARE-ACQUISITION-TO-MAXIMIZE-LETHALITY.PDF

Hegseth, P. (2025, November 7). *Memorandum on Transforming the Defense Acquisition System into the Warfighting Acquisition System to Accelerate Fielding of Urgently Needed Capabilities to Our Warriors*. https://media.defense.gov/2025/Nov/10/2003819439/-1/-1/1/TRANSFORMING-THE-DEFENSE-ACQUISITION-SYSTEM-INTO-THE-WARFIGHTING-ACQUISITION-SYSTEM-TO-ACCELERATE-FIELDING-OF-URGENTLY-NEEDED-CAPABILITIES-TO-OUR-WARRIORS.PDF

Hegseth, P. (2026, January 9). *Memorandum on the Artificial Intelligence Strategy for the Department of War*. https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/ARTIFICIAL-INTELLIGENCE-STRATEGY-FOR-THE-DEPARTMENT-OF-WAR.PDF

H.R. 1588. National Defense Authorization Act for Fiscal Year 2004. 108th Congress. (2003). https://www.congress.gov/bill/108th-congress/house-bill/1588/text

H.R. 2670. National Defense Authorization Act for Fiscal Year 2024. 118th Congress. (2023). https://www.congress.gov/bill/118th-congress/house-bill/2670/text

H.R. 7776. James M. Inhofe National Defense Authorization Act for Fiscal Year 2023. (2022). https://www.congress.gov/bill/117th-congress/house-bill/7776

Institute for Ethics in Government. (2021, March 18). *Ethics Laws Applicable to Special Government Employees*. https://extapps2.oge.gov/Training/OGETraining.nsf/%24%24OpenDominoDocument.xsp?documentId=D006291C1FEC02448525869C005BD4B8&action=openDocument

Kaplan, J., McCandlish, S., et al. (2020, January 23). Scaling Laws for Neural Language Models. *ArXiv*. https://arxiv.org/abs/2001.08361

Kingston, S. K. (2026, March 23). *State Department Launches Effort to Counter Cyberattacks, AI Risks from Iran, Others*. ABC News. https://abcnews.com/Politics/state-department-launches-effort-counter-cyberattacks-ai-risks/story?id=131265350

Kupor, S. (2025, December 15). *OPM Memorandum on Building the AI Workforce of the Future.* Office of Personnel Management. https://www.opm.gov/chcoc/latest-memos/building-the-ai-workforce-of-the-future.pdf

LeCun, Y., Bengio, Y., & Hinton, G. (2015, May 28). Deep learning. *Nature 521*. https://www.cs.toronto.edu/~hinton/absps/NatureDeepReview.pdf

Levels.fyi. (n.d.). *Anthropic salaries.* Retrieved March 10, 2026, from https://www.levels.fyi/companies/anthropic/salaries

Lutnick, H. (2025, June 3). *Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation.* U.S. Department of Commerce Press Release. https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai

Mazur, J. (n.d.). *The Federal Government Needs to Expand Its Technology Talent Pipelines*. Retrieved March 10, 2026, from https://www.careersingovernment.com/tools/gov-talk/career-advice/the-federal-government-needs-to-expand-its-technology-talent-pipelines/

McQuade, J. & Murray, M., et al. (2019, May 3). *Software Is Never Done: Refactoring the Acquisition Code for Competitive Advantage*. Defense Innovation Board. https://media.defense.gov/2019/May/01/2002126690/-1/-1/0/SWAP%20EXECUTIVE%20SUMMARY.PDF

Nevo, S. (2025, August 11). *A Sprint Toward Security Level 5*. Institute for Progress. https://ifp.org/a-sprint-toward-security-level-5

Obis, A. (2025, March 19). *CENTCOM scales its AI infrastructure, shaping future of AI for combatant commands*. Federal News Network. https://federalnewsnetwork.com/artificial-intelligence/2025/03/centcom-scales-its-ai-infrastructure-shaping-future-of-ai-for-combatant-commands

OpenAI. (2025, September 12). *Working with US CAISI and UK AISI to build more secure AI systems*. https://openai.com/index/us-caisi-uk-aisi-ai-update/

OpenAI. (2026). *Economic Espionage Complaint*. https://assets.bwbx.io/documents/users/iqjWHBFdfxIU/rRmql_jJcxb4/v0

Owen, D. (2025, September 16). *AI in 2030: Extrapolating Current Trends*. Epoch AI. https://epoch.ai/files/AI_2030.pdf

Pendleton, J. H. (2010, June 9). *National Security: Key Challenges and Solutions to Strengthen Interagency Collaboration (GAO-10-822T)*. U.S. Government Accountability Office. https://www.gao.gov/assets/gao-10-822t.pdf

Pilz, K., Rahman, R., Sanders, J., & Heim, L. (2026, March). *Data on GPU Clusters*. Epoch AI. https://epoch.ai/data/gpu-clusters

Press, R. (2026). *CAISI Issues Request for Information About Securing AI Agent Systems*. National Institute for Standards & Technology. https://www.nist.gov/news-events/news/2026/01/caisi-issues-request-information-about-securing-ai-agent-systems

Reddy, P., & Gujral, A. (2025). *EchoLeak: The First Real-World Zero-Click Prompt Injection Exploit in a Production LLM System*. https://arxiv.org/pdf/2509.10540

Rubio, M. (2025). *Next Steps on Building an America First State Department*. https://www.state.gov/releases/office-of-the-spokesperson/2025/05/next-steps-on-building-an-america-first-state-department

S.2551. Artificial Intelligence Training for the Acquisition Workforce Act. 117th Congress. (2022). https://www.congress.gov/bill/117th-congress/senate-bill/2551

Serbu, J. (2023, November 3). *Navy says two programs show the case for 'colorless' IT spending*. Federal News Network. https://federalnewsnetwork.com/navy/2023/11/navy-says-two-programs-show-the-case-for-colorless-it-spending/.

Sevilla, J., et al. (2022, March 9*). Compute Trends Across Three Eras of Machine Learning*. https://arxiv.org/pdf/2202.05924

Shibu, S. (2025, August 8). *Mark Zuckerberg Reportedly Made One Person a $1.5 Billion Job Offer — and Was Rejected. Here's How Google, Microsoft, and OpenAI Are Competing with Meta in the AI Talent Wars*. Entrepreneur. https://www.entrepreneur.com/business-news/meta-makes-billion-dollar-job-offer-competing-for-ai-talent/495672

Shriver, R. (2025). *Federal Employee Viewpoint Survey Results*. https://www.opm.gov/fevs/reports/governmentwide-reports/governmentwide-reports/governmentwide-management-report/2024/2024-governmentwide-management-report.pdf

Singapore Infocom Media Development Authority. (2024, May 22). *Digital Trust Centre designated as Singapore's AI Safety Institute*. https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2024/digital-trust-centre

Somala, V. & Cottier, B. (2025, November 10). *Build times for gigawatt-scale data centers can be 2 years or less*. Epoch AI. https://epoch.ai/data-insights/data-centers-buildout-speeds

The White House. (2025a, April 18). *Fact Sheet: President Donald J. Trump Creates New Federal Employee Category to Enhance Accountability*. https://www.whitehouse.gov/fact-sheets/2025/04/fact-sheet-president-donald-j-trump-creates-new-federal-employee-category-to-enhance-accountability

The White House. (2025b, July 24). *Winning the Race: AMERICA'S AI ACTION PLAN.* https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf

The White House. (2025c, September 18). *Memorandum of Understanding between the Government of the United States Of America and the Government of the United Kingdom of Great Britain and Northern Ireland regarding the Technology Prosperity Deal.* https://www.whitehouse.gov/presidential-actions/2025/09/memorandum-of-understanding-between-the-government-of-the-united-states-of-america-and-the-government-of-the-united-kingdom-of-great-britain-and-northern-ireland-regarding-the-technology-prosperity-de

UK Department for Science, Innovation, & Technology. (2026, January 29). *AI Opportunities Action Plan: One Year On*. https://www.gov.uk/government/publications/ai-opportunities-action-plan-one-year-on/ai-opportunities-action-plan-one-year-on

U.S. Center for AI Standards and Innovation (CAISI). (2025). *Strengthening AI Agent Hijacking Evaluations*. NIST Technical Blog. https://www.nist.gov/news-events/news/2025/01/technical-blog-strengthening-ai-agent-hijacking-evaluations

U.S. Department of Justice: Justice Management Division. (2006, February 6). *Summary of Government Ethics Rules for Special Government Employees.* https://www.justice.gov/jmd/ethics/summary-government-ethics-rules-special-government-employees

U.S. Department of Energy. (2025, October 28). *Energy Department Announces New Partnership with NVIDIA and Oracle to Build Largest DOE AI Supercomputer*. https://www.energy.gov/articles/energy-department-announces-new-partnership-nvidia-and-oracle-build-largest-doe-ai

U.S. Department of War. (2026). *War Department Launches AI Acceleration Strategy to Secure American Military AI Dominance*. https://www.war.gov/News/Releases/Release/Article/4376420/war-department-launches-ai-acceleration-strategy-to-secure-american-military-ai/

U.S. General Services Administration. (2024, December 4). *Empowering responsible AI: How expanded AI training is preparing the government workforce. GSA Blog*. https://www.gsa.gov/blog/2024/12/04/empowering-responsible-ai-how-expanded-ai-training-is-preparing-the-government-workforce

U.S. Government Accountability Office. (2025a). *GAO-25-107653: Generative AI Use and Management at Federal Agencies.* https://www.gao.gov/assets/gao-25-107653.pdf

U.S. Government Accountability Office. (2025b). *GAO-25-108519: Science & Tech Spotlight: AI Agents*. https://www.gao.gov/assets/gao-25-108519.pdf

USA Jobs. (2026, February 3). *Job Posting: AI Standards Architect*. https://www.usajobs.gov/job/856268900

USA Jobs. (2026, February 3). *Job Posting: Senior Cyber Offense Specialist*. https://www.usajobs.gov/job/856268500

U.S. National Institute for Standards & Technology (NIST). *CAISI Research Blog: A NIST blog from the Center for AI Standards and Innovation*. Retrieved [March 20, 2026] from https://www.nist.gov/blogs/caisi-research-blog

U.S. National Institute for Standards & Technology (NIST). (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

U.S. Office of Personnel Management. (2025a). *Salary Table 2025-DCB*. https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/pdf/2025/DCB.pdf

U.S. Office of Personnel Management. (2025b). *Salary Table No. 2025-EX*. https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/pdf/2025/EX.pdf

U.S. Office of Personnel Management. (2026a). *Fact Sheet: Aggregate Limitation on Pay*. https://www.opm.gov/policy-data-oversight/pay-leave/pay-administration/fact-sheets/aggregate-limitation-on-pay

U.S. Office of Personnel Management. (2026b). *Intergovernment Personnel Act*. https://www.opm.gov/policy-data-oversight/hiring-information/intergovernment-personnel-act

U.S. Senate Committee on Appropriations. (2024, March 3). *BILL SUMMARY: Commerce, Justice, Science, and Related Agencies Fiscal Year 2024 Appropriations Bill*. https://www.appropriations.senate.gov/imo/media/doc/fy24_cjs_bill_summary.pdf

U.S. Tech Force. (n.d.). *Tech for the American People.* Retrieved March 20, 2026, from https://techforce.gov

Vought, R. (2025, April 4). *OMB Memorandum M-25-21: Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*. https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf

Waterman, S. (2025, May 1). *Meeting the Software Challenge: How the New Acquisition Pathway Came to Be*. Air and Space Forces Magazine. https://www.airandspaceforces.com/software-acquisition-pathway-history-part-2

Author Biographies

Cole Salvador *is a Fellow for AI and Emerging Technology at the America First Policy Institute, where he focuses on issues like data centers and energy, AI and national security, and international competition.*

Jack Crovitz *is a Senior Fellow for AI and Emerging Technology at the America First Policy Institute, focused on AI's role in cybersecurity and American national security dominance. He holds a B.A. from the University of Chicago, where he studied economics and law.*

Yusuf Mahmood *is Director of AI and Emerging Technology at AFPI. He holds a J.D. from Harvard Law School and two B.A.s in economics and philosophy from the University of Maryland, College Park.*