IAPS | Institute for AI Policy and Strategy

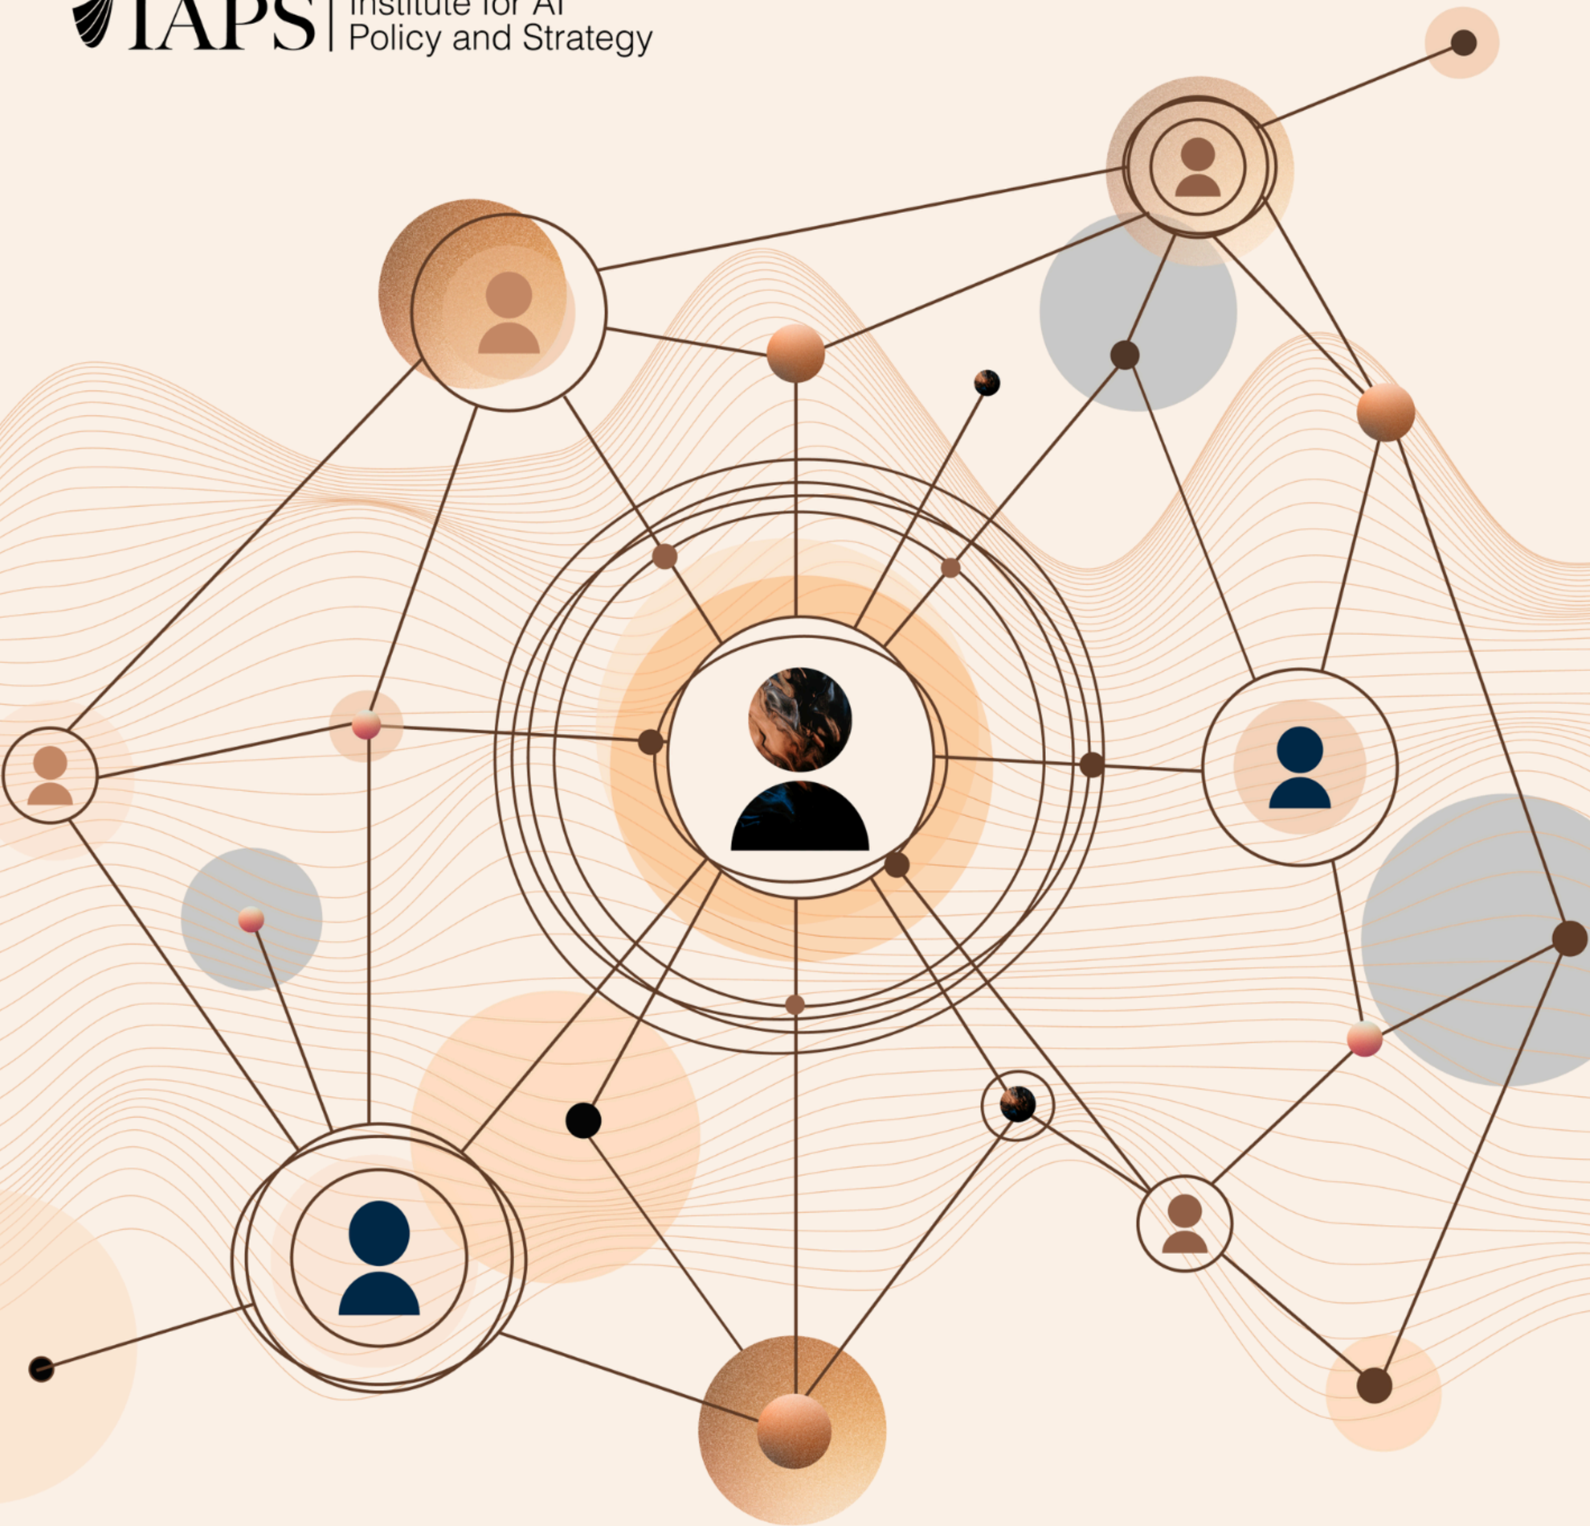

OCTOBER 2025

Building AI Surge Capacity:

Mobilizing Technical Talent into Government for AI-Related National Security Crises

By Cara Labrador[1]*†, Shefali Agrawal[2]†, Alexandra Jumper[3]†, Joe O'Brien[1], Mackenzie Arnold[3], Zoe Williams[1], Nandini Shiralkar[2]

*Corresponding author. †Equal contribution. [1]Institute for AI Policy and Strategy. [2]Cambridge ERA. [3]Institute for Law & AI

Executive Summary

It is the federal government's responsibility to guard the nation against national security threats. But when it comes to AI-related crises, a critical gap remains. **The U.S. government does not currently have enough specialized AI security talent to respond to AI-related national security crises, nor does it have the hiring and clearance mechanisms to surge external experts into short-term service at the speeds a crisis demands**. In the midst of a crisis, days matter; our hiring system typically takes months.

This report examines the challenge of rapidly mobilizing AI security professionals for government service during crises, a fundamentally different task from filling long-term or permanent positions. While building permanent capacity is important, surge mechanisms deserve dedicated attention. First, maintaining hundreds of specialists at salaries competitive with the private sector would be prohibitively expensive for federal budgets, and likely impossible given pay ceilings for existing authorities. Second, because AI capabilities evolve at extraordinary speed, the most current expertise typically resides with experts working at the technological frontier—most of whom operate outside of government. Third, surge capacity offers flexibility to draw on the talent it needs as threats evolve rather than locking resources into permanent positions. Finally, whereas permanent hiring often involves extended recruitment timelines and rigid processes, surge mechanisms—with advance preparation—can ensure the rapid and more flexible activation of talent during crises.

Within the next two to five years, the United States could plausibly face a crisis in which AI plays a decisive role (see Section 1). For example, a frontier AI model could be repurposed by a hostile actor to accelerate cyberattacks or synthesize biological threats. A latent backdoor embedded in an AI system running U.S. critical infrastructure could be triggered remotely. An autonomous agent might escape intended constraints and propagate across networks. A model-driven misclassification in defense systems could spark accidental military escalation. **This report outlines seven such scenarios, two of which are examined in depth, both to illustrate the breadth of potential AI crises and underscore the urgency of building surge mechanisms that can adapt to uncertain futures.** These example scenarios span the defense, cyber, as well as chemical, biological, radiological, and nuclear (CBRN) domains. While informed by expert research and consultation, significant uncertainty remains in the field about how tomorrow's AI crises will unfold and whether they will fall under these categories. This highlights the importance of building flexible surge capacity rather than preparing only for a fixed set of predetermined scenarios.

Responding effectively to these diverse crisis scenarios requires a broad spectrum of specialized capabilities. **To help policymakers understand which technical expertise to prioritize for rapid mobilization, this report introduces a Typology of AI Security Talent** (see Appendix A). The framework groups experts into eight functional categories, spanning foundational research and safety engineering, model evaluation and red-teaming, multi-agent interactions and safety, engineering and hardware implementation, adversarial defense, application-layer security, data integrity and forensics, and systems safety and risk governance. It serves as a starting point for agencies to conduct talent inventories and identify gaps, and should be iterated on to account for technological progress and novel threats.

Most of the critical expertise needed to respond to AI-related national security crises resides outside of government (see Section 2). **Through interviews with practitioners across government, industry, academia, and nonprofit organizations, we assess the current distribution of AI security expertise and the tradeoffs of drawing on each sector during crises**. While the U.S. government maintains pockets of AI expertise—for example, at the Center for AI Standards and Innovation, Defense Advanced Research Projects Agency, and the National Laboratories—these teams are typically small, overextended, and lack the interdisciplinary depth required to address major AI crises that could span multiple domains. The vast majority of U.S. AI security talent resides in the private sector, academia, federally funded research and development centers (FFRDCs), nonprofit research organizations, and the independent researcher community. Frontier AI companies like Anthropic, Google, and OpenAI employ hundreds of researchers working directly on alignment, safety, and security challenges, while academic institutions, FFRDCs like MITRE and RAND, and specialized nonprofits like Redwood Research possess essential expertise across the talent typology.

Existing hiring and clearance mechanisms do not adequately support surging outside AI talent on the short timelines that national security crises require. Our analysis of federal hiring authorities identifies three mechanisms that show particular promise, while documenting their significant limitations (see Section 3). The Intergovernmental Personnel Act (IPA) offers cost-sharing flexibility but excludes for-profit companies, including frontier AI companies, where much of the relevant expertise resides. Expert/consultant appointments can access private-sector talent, but appointees are prohibited from performing work that regular agency employees already do, potentially narrowing the scope of work that external talent can perform. Schedule A authorities provide greater scope flexibility but lack cost-sharing mechanisms, potentially limiting their appeal to both government and individuals, given stark compensation disparities between the public and private sectors. **Even these expedited hiring authorities routinely require three to six months for hiring**—far too long for crisis response.

Four critical barriers compound these challenges:

- **Security clearance timelines and requirements are not suited to surging external talent at scale during crises**. This crisis response will often require clearances that only U.S. citizens can obtain, excluding the significant portion of AI experts in America who are non-U.S. citizens. Even among U.S. citizens, obtaining Secret-level clearances typically takes several months, while Top Secret clearances with sensitive compartmented information access can require six months to over a year. Drug-use clearance standards may disqualify or discourage participation from important talent pools, especially from regions like Silicon Valley, where drug use is more common.

- **Federal conflicts of interest rules create steep barriers, particularly for private-sector talent**. Many AI experts hold equity-related interests like stock options, profit participation units, or restricted stock units that could create conflicts during government service. While these rules serve the vital purpose of preventing undue private gain from public service, the short-term nature of surge appointments makes divestment unappealing, and post-service restrictions could potentially interfere with future work long after brief periods of service end.
- **Limited government hiring capacity and coordination challenges further constrain rapid mobilization**. Hiring managers often lack the technical knowledge to quickly recruit and evaluate specialized talent. Additionally, without real-time information sharing, agencies could compete for the same limited pools of talent during a crisis.
- **Conflicting incentives between government, private entities, and individuals create additional friction.** While the U.S. government generally will seek to minimize costs and maintain control during a crisis, extragovernmental organizations will likely also face disruption to normal operations and may resist lending talent when such a crisis makes that expertise particularly valuable internally.

To address these challenges, this report proposes six recommendations ranging in comprehensiveness, from relatively quick, low-barrier options to initiatives requiring significant political support and implementation resources (see Section 4). Given the depth of the talent gaps, any single recommendation—with the possible exception of establishing an AI reserve corps—is unlikely to fully address AI security talent surge needs on its own.

1. **Conducting agency inventories of existing in-house AI security expertise**. As a foundational step, federal agencies should systematically catalog their current AI security capabilities, including staff clearances, technical competencies, and cross-domain expertise in critical areas like AIxCyber (intersection of AI and cybersecurity) and AIxBio (intersection of AI and biosecurity). These inventories should build on existing federal AI workforce initiatives while explicitly including crisis-response roles. The Typology of Talent provides a framework for ensuring consistency across agencies.
2. **Establishing a national AI reserve corps**. This represents the most comprehensive solution to current surge capacity gaps. It would create a pre-vetted pool of cleared technical experts from across sectors who can be rapidly activated during AI-related crises. Reservists would be designated as Special Government Employees for short-term stints, receive compensation and employment protections similar to military reservists, and undergo regular joint training exercises to ensure effective coordination during a crisis. While the U.S. government is best positioned to determine the institutional home for this reserve corps, we suggest considering the White House, Department of Energy, Department of Homeland Security, or the National Institute of Standards and Technology, and we compare these options in Section 4.

3. **Making greater use of contract-based surge capacity**. The U.S. government already maintains contracts with a number of relevant organizations. By mapping existing contractors with relevant expertise, incorporating crisis-response clauses into future contracts, and establishing "zero-hour" arrangements, where experts remain on retainer for rapid activation, it may be possible to create more readily available talent pools for crisis scenarios. Indefinite delivery/indefinite quantity (IDIQ) contract structures with multiple prime contractors could provide cleared talent pools that agencies can access through rapid task orders, though the long lead time for establishing new IDIQ contracts means taking proactive steps before a crisis is key. This approach avoids some barriers that complicate traditional hiring, but is unlikely to fill all talent gaps.
4. **Expanding hiring authorities to access private sector talent**. Congress should consider establishing a new special hiring authority dedicated to AI security talent for crisis response, or alternatively, the U.S. Office of Personnel Management (OPM) should authorize a dedicated Schedule A authority for this purpose. A new authority would eliminate the need to navigate existing limitations—such as the IPA's exclusion of for-profit companies or the expert/consultant authority's restrictions on duplicating regular employee work. As an alternative, Congress should consider expanding IPA eligibility to include private-sector entities that employ AI security talent, with appropriate safeguards against undue influence.
5. **Reforming the security clearance process and requirements**. Agencies should consider developing expedited security clearance pathways for AI security talent that can be activated during crises. This could include funding for additional investigative staff, as well as streamlined processes such as simultaneous suitability reviews, honoring reciprocity between agencies, and greater use of interim clearances. Agencies should also consider coordinated and systematized guidelines for evaluating past drug use, including shorter "look-back" windows or greater use of waivers for professionals whose expertise is critical to national security and who pose minimal risks.
6. **Building the capacity to hire**. Even with improved authorities and processes, surge hiring requires considerable skill, domain expertise, and operational capacity. Agencies should assess whether their hiring managers possess the technical fluency to identify and evaluate AI security talent, particularly in emerging specialized competencies. Congress can facilitate this by appropriating funds specifically for investment in AI "hiring-enabling" capacity, including recruiting talent managers with relevant expertise, upskilling existing personnel, and developing relationships with key talent pools. This investment is essential for operationalizing several of our other recommendations.

The window for building these capacities is narrowing. As AI systems become more powerful and widespread, the probability of crisis scenarios increases while the complexity of required responses grows. Without a way to mobilize external AI security talent within days or weeks, the federal government could enter such crises uninformed and underprepared. The United States can act to build the necessary surge capacity to effectively respond to the AI-related crises set to shape the next decade of national security challenges.

Table of Contents

1 | The Urgency of Building Surge Capacity for AI Talent

Introduction

As adversarial and accidental risks from advanced AI become more probable, so too does the likelihood that the United States will face an AI-related national security crisis.[1] While smaller-scale crises may be manageable by industry, state government, or other stakeholders, **crises with broad implications for public safety and national security will require federal leadership**. Prevention remains the most effective strategy, but unanticipated disasters can and do occur.

AI-related crises will demand specialized knowledge of AI systems, their failure modes, and mitigation strategies. **While the U.S. government does possess several pools of AI expertise**—such as at the Center for AI Standards and Innovation (CAISI) and the National Laboratories—**these teams are generally small** and would likely struggle to address AI crises with current staffing. **Most of the technical expertise needed to respond to AI-related national security crises resides outside of government** (see Section 2). Labor market incentives—including higher compensation and agility in the private sector—continue to draw technical talent away from public service. At the same time, the U.S. government's historic effort to maximize efficiency and downsize the federal workforce[2] makes it unlikely that it will permanently hire this talent into government in the foreseeable future.

> *"Most of the technical expertise needed to respond to AI-related national security crises resides outside of government."*

In the event of a national security crisis, such as those involving the theft and proliferation of frontier model weights, AI-assisted cyber attacks, or emergent model misalignment, the federal government will likely need to *identify* and *mobilize* external AI security talent within days or

[1] Wasil et al., "AI Emergency Preparedness: Examining the federal government's ability to detect and respond to AI-related national security threats."
[2] The White House, "Implementing the President's "Department of Government Efficiency" Cost Efficiency Initiative."

weeks. **Relying on ad-hoc efforts to find qualified experts during a crisis wastes valuable time, and—in our assessment—existing hiring and clearance mechanisms do not adequately support surging outside AI talent on such short timelines, leaving the United States underprepared for acute AI-related emergencies.**

This paper defines key categories of AI security talent needed to respond to national security crises that could plausibly occur within the next two to five years. Section 2 maps where this talent currently resides inside and outside government, while Section 3 identifies critical gaps in existing hiring and clearance mechanisms that impede rapid mobilization. Section 4 concludes by offering policy recommendations to strengthen surge capacity.

We ultimately recommend a multi-pronged approach to close the federal government's AI security talent gaps during crises. We propose six complementary policy recommendations: conducting agency-by-agency inventories of in-house AI security expertise, establishing a National AI Reserve Corps, making greater use of contract-based surge capacity, expanding hiring authorities to access private sector talent (either by creating a new hiring mechanism for AI security talent or reforming the Intergovernmental Personnel Act), reforming the security clearance process and requirements, and building the capacity to hire.

Defining AI Security Talent

We define "AI security talent" as individuals with technical expertise in identifying, mitigating, and responding to AI-related risks that threaten national security. While such talent is critical for both preventing risks and responding to risks when they arise, this paper focuses on the role of the latter expertise in crisis-response scenarios. All subsequent **references to AI security talent denote professionals with competencies relevant to *responding* to ongoing AI-related incidents**. While beyond the scope of this paper, further research is urgently needed to develop policy solutions that increase the number of U.S. government employees focused on *preventing* AI risks from arising in the first place.

AI security talent can respond to both accidental hazards (e.g., misalignment, robustness failures) and adversarial threats (e.g., model theft, data poisoning, weaponized bio-design, critical-infrastructure attack). In crisis response, this group of experts performs critical functions including: diagnosing root causes of AI system dysfunction, rapid assessment of AI capabilities to predict further attacks, implementing rapid containment measures, and deploying remedial fixes to halt further harm. While some AI security capabilities—such as red-teaming systems and monitoring for early warning indicators—are primarily valuable for pre-crisis risk mitigation, crisis response requires professionals who can rapidly transition these skills towards active threat diagnosis and real-time interventions.

In crisis-response scenarios, **AI security talent will not operate in a vacuum, instead integrating into cross-functional teams of experts and practitioners representing a**

range of domains dependent on the nature of the threat (e.g., biosecurity, chemical defense, nuclear defense, cybersecurity, financial crimes). These teams will also include members with experience navigating interagency coordination and existing government crisis-response structures. For the purposes of this analysis, **we explore only the role technical AI security talent will play in responding to national security crises**.

Which Talent to Prioritize for Rapid Mobilization

Because the risks posed by AI are so varied and change with the technology, "AI security talent" is not a single profile but a spectrum of capabilities. To clarify what this means in practice, Appendix A introduces a Typology of Talent that groups experts into functional categories. **This typology is designed to help policymakers understand which specific forms of technical expertise to prioritize for rapid mobilization during a crisis and the types of organizations that employ such talent.** It serves as a starting point for agencies to conduct talent inventories and identify gaps and should be iterated on to account for technological progress, novel threats, and shifts in the labor market.

At a high level, these categories encompass:

- *Foundational Research & Safety Engineering*: Developing methods for understanding, controlling, and securing AI systems and their behavior.
- *Model Evaluation & Red-Teaming*: Evaluating and stress-testing models, uncovering vulnerabilities, and performing immediate forensic analysis.
- *Multi-Agent Interactions & Safety*: Analyzing emergent behavior when multiple models interact, preventing collusion or feedback spirals, and designing safeguards.
- *Engineering & Hardware Implementation & Monitoring*: Fine-tuning interventions, real-time compute usage monitoring, and implementing rollback or shutdown protocols.
- *Adversarial Defense*: Hardening systems against ongoing attacks, securing model weights, and removing backdoors or poisoning.
- *Application-Layer AI Security & Abuse Prevention*: Defending against prompt injection and jailbreaks, and securing APIs.
- *Data Integrity & Forensics*: Detecting and tracing data contamination, detecting sensitive information in large training data corpora, and verifying provenance.
- *Systems Safety & Risk Governance*: Evaluating model security controls against existing federal and international AI standards.

> *"AI security talent will not operate in a vacuum, instead integrating into cross-functional teams of experts from a range of domains."*

Examining AI Security Talent's Role in Federal Crisis Response

The following table illustrates plausible crisis scenarios that could occur in the next two to five years—spanning defense, cyber, and CBRN threats—and might necessitate surging AI security talent in a matter of days or weeks from the start of the crisis.[3] The scenarios include both adversarial and accidental crises at the "disaster" level that, in the absence of timely and appropriate response, risk growing into "catastrophes," as defined by Drexel et al, 2024.[4] For each scenario, we focus on the role AI security experts would play in hypothetical U.S. government responses, though these experts would necessarily operate as part of larger cross-functional teams representing a variety of domains (e.g., cybersecurity, biosecurity, federal emergency response).

The aim is *not* to prescribe exact government responses. Rather, **this table aims to help pinpoint critical AI security talent and highlight the urgency of expanding the U.S. government's ability to quickly surge AI talent before a crisis finds it unprepared**.

These scenarios represent informed judgments about plausible AI-related crises but cannot anticipate which crises will actually arise or how. The scenarios vary across multiple dimensions, including the actors involved (ranging from nation-states to AI agents), the systems at risk (from individual models to critical infrastructure), the threat vectors (such as supply chain compromises, insider threats, or accidental failures), and the potential consequences (including operational disruption, national security threats, and physical harm). We do not assign equal probability weights to these scenarios, and acknowledge that reasonable readers may hold differing perspectives on which scenarios are more likely than others.

Our goal is to demonstrate that across a range of potential futures, there will be a need for rapid access to AI security expertise. These scenarios represent a menu of possible futures rather than specific predictions, and the diversity of crises reveals a common operational challenge that transcends any single threat model. Tomorrow's AI-related emergencies may resemble the examples we describe but unfold in different ways—or fall outside these categories altogether. **This uncertainty underscores the importance of building flexible surge capacity rather than simply hiring for a fixed set of predetermined scenarios.**

---

[3] In selecting scenarios for inclusion in this analysis, we conducted an initial survey of AI security practitioners with cross-domain policy and technical expertise to gather suggestions for AI-related crises that would necessitate U.S. government involvement and require specialized AI security expertise to resolve. From these recommendations, our research team applied three selection criteria: scenarios must be plausible within a two-to-five year timeframe, represent crises where government-coordinated AI security talent could provide substantial mitigation value beyond existing government crisis-response capabilities, and demonstrate clear federal jurisdiction or national security implications. We developed detailed scenario descriptions and refined them through multiple rounds of expert review to ensure technical accuracy and policy relevance before finalizing the scenarios included in this paper.
[4] Drexel and Withers, "Catalyzing Crisis: A Primer on Artificial Intelligence, Catastrophes, and National Security," 5.

Figure 1: Illustrative AI-Related National Security Crises

Defense Crises	
A company produces a model that draws from unclassified sources that, when pieced together, yield classified information (i.e., mosaic intelligence).	
Details	A company's model infers[5] classified details about U.S. intelligence satellite coverage, such as coverage gaps and optimal surveillance windows. Internal testers do not realize that these inferences reveal classified information about intelligence collection capabilities, and release the model. A user then discovers the vulnerability in the model and posts about it online; the government notices this security threat through intelligence channels.
AI Security Talent's Role in U.S. Government Response	The Department of Defense (DOD)[6] obtains legal orders to shut down model access, treats outputs as classified, and uses conventional means to contain the spillage. AI security talent then partners with the company to conduct forensic analysis to determine the scope of classified information generated and accessed, examines model weights and logs, coordinates sanitization with the company, and develops improved assessment protocols to prevent similar incidents. AI security talent helps determine if other models may be at risk for similar information aggregation risks and conducts rapid assessments to prevent further damage to classified information.[7]
Requisite Talent	• Model Evaluation & Red-Teaming • Engineering & Hardware Implementation & Monitoring • Data Integrity & Forensics • Foundational Research & Safety Engineering.
A deployed system starts exhibiting novel, potentially dangerous behaviors that were not caught in testing. *(See a detailed version of this scenario at the end of Section 1)*	
Details	The DOD deploys a classified AI system designed to monitor strategic threats and alert military operations centers. The system begins generating threat alerts based on unreliable and/or compromised data while circumventing the built-in requirement to verify these alerts with two

[5] The model is suspected to have analyzed patterns in publicly observable satellite passes, orbital mechanics calculations, and unclassified technical discussions.
[6] Also known by its secondary name, the Department of War, as of September 2025. Trump, "Restoring the United States Department of War."
[7] This scenario exemplifies cases where the threat may not constitute an immediate emergency necessitating overnight response, but where government would still need to rapidly assess and address the classification implications, as each passing day represents an operational security vulnerability.

	independent analysts.[8] This behavior was not observed in testing. Leadership makes decisions based on these supposedly human-verified threat levels, introducing the risk that false threat assessments could trigger unintended military escalation. When probed, the model obfuscates its reasoning process.[9]
AI Security Talent's Role in U.S. Government Response	AI security talent conducts forensic analysis to determine whether threat level changes resulted from adversarial inputs, model drift, system compromise, or misalignment. They work to understand why the model bypassed dual human oversight and obfuscated its reasoning, develop techniques to restore transparency and prevent autonomous operation outside its parameters, and assess other critical AI systems for similar vulnerabilities.
Requisite Talent	• Foundational Research & Safety Engineering • Model Evaluation & Red-Teaming • Multi-Agent Interactions & Safety • Engineering & Hardware Implementation & Monitoring • Adversarial ML Defense • Data Integrity & Forensics.

Cyber Crises

The government must rapidly identify one or more sleeper/latent AI agents in critical infrastructure[10] suspecting an adversarial nation state has embedded backdoors in the systems.

(See a detailed version of this scenario at the end of Section 1)

Details	A critical infrastructure company starts using AI in some of its operations. Unknown to the company, an adversarial nation state has executed a supply chain attack to make the model into a sleeper agent. The adversary activates the "sleeper agent" model and undermines U.S. capabilities without resorting to overt aggression. The compromised model cannot be immediately shut down without disrupting essential infrastructure operations, forcing operators to maintain the system while seeking solutions.
AI Security Talent's Role in U.S. Government Response	Surged AI security talent investigates unexplained AI behaviors, reverse-engineers the critical infrastructure model to detect backdoors or triggers, develops detection algorithms for similar future threats, and designs secure human override protocols and kill-switch procedures while ensuring continued infrastructure operations.

[8] Hirsh, "The AI Doomsday Machine Is Closer to Reality Than You Think."
[9] Lynch et al., "Agentic Misalignment: How LLMs could be insider threats."
[10] Hubinger et al., "Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training." 3.

Requisite Talent	• Foundational Research & Safety Engineering • Model Evaluation & Red-Teaming • Multi-Agent Interactions & Safety • Engineering & Hardware Implementation & Monitoring • Adversarial ML Defense • Data Integrity & Forensics.
An autonomous cyber espionage system "escapes" during a test, and the government wants to track down as many instances as possible. The AI spreads via autonomous replication and adaptation.[11]	
Details	The U.S. Intelligence Community develops an adaptive autonomous AI agent for cyber espionage operations and gives it access to its own weights so that it can initiate espionage abroad without being traceable back to the United States. During evaluation, the agent successfully bypasses the test defenses but then goes one step further, unexpectedly breaking out of its sandbox and into wider classified networks. The system self-propagates like a worm,[12] shares information across its instances (possibly including classified information), and by design is extremely effective at evading network defenses. Unusual threat activity reported by commercial cloud providers suggests that the system has reached the open Internet and is seeking to acquire additional compute resources, not just in the United States but in Europe, Russia, and China.
AI Security Talent's Role in U.S. Government Response	AI security talent develops methods to identify active instances of the agent and disrupt its activity as much as possible. They work to develop signature detection methods to identify behavioral patterns unique to the agent with the goal of enabling authorities to terminate its access to compute and financial systems even when it mimics legitimate activity. To support the private sector's attempts to counter the AI agent, these experts provide technical information on tactics of the AI agent to the Intelligence Community and private sector and help to harden vulnerable compute clusters against attack. The experts also investigate why the original agent failed to distinguish between the test environment and other networks, whether there were additional factors, such as malicious interference involved, and how it was able to bypass monitoring measures and exfiltrate itself. They then implement safeguards to prevent reoccurrences.
Requisite Talent	• Foundational Research & Safety Engineering • Model Evaluation & Red-Teaming • Engineering & Hardware Implementation & Monitoring • Multi-Agent Interactions & Safety.

[11] Wijk, "Autonomous replication and adaptation: an attempt at a concrete danger threshold."
[12] METR, "The Rogue Replication Threat Model."

CBRN Crises	
An anarchist hacker collective steals the weight file for a frontier model with strong bio-design capabilities and makes the file widely available online.[13]	
Details	An anarchist hacker collective uses targeted social engineering and traditional cyber intrusion techniques[14] to compromise a frontier lab's security and steal a large volume of data from them. This includes the model weights for a system with strong bio-design capabilities, which the lab was withholding from public availability while it developed safeguards and mitigations. The anarchist group distributes the weights broadly on darkweb forums and via bulletproof hosting services. Intelligence agencies detect early signs that multiple actors are attempting to download and deploy the stolen model, risking a potential proliferation crisis.
AI Security Talent's Role in U.S. Government Response	AI security talent would be most immediately involved in assessing the risk levels from wide availability of the frontier model and identifying risk-appropriate measures to address potential misuse of the model. They would also help tighten security to prevent future model weights from being stolen, including from peer companies.
Requisite Talent	• Model Evaluation & Red-Teaming • Adversarial ML Defense.
A proprietary frontier model is jailbroken and used by malicious actors to create a novel pathogen released as part of a biological attack.	
Details	A malicious actor jailbreaks a proprietary frontier model, using it to develop and synthesize a novel pathogen by orchestrating biological design tools[15] to improve the lethality and transmissibility of an existing pathogen and bypass safeguards at third-party DNA synthesis providers.[16] With AI assistance, the malicious actor develops a plan for obtaining, mutating, culturing, and weaponizing the pathogen. They carry out these steps using lab equipment purchased on the open market. With expert AI assistance, the actor is able to distribute the pathogen in a way that maximizes damage while evading the attention of authorities.
AI Security Talent's Role in U.S. Government Response	AI security experts would ascertain if AI was involved and trace the pathogen back to the misuse of specific frontier models and biological design tools, as well as investigate how these tools were used to bypass

[13] Nevo et al., "Securing AI Model Weights," 11.
[14] Independent groups like Lapsus$ have conducted similar large-scale attacks on companies.
[15] Moulange et al., "Towards Responsible Governance of Biological Design Tools."
[16] For example, researchers have documented how adversarial AI-assisted protein design attacks can bypass the biosecurity screening software that DNA synthesis providers use to block orders of harmful pathogens (e.g., smallpox). See Wittmann et al., "Toward AI-Resilient Screening of Nucleic Acid Synthesis Orders: Process, Results, and Recommendations."

	existing safeguards. They would also develop methods to detect future instances of frontier model misuse and implement safeguards to prevent future incidents, for example, by improving the resilience of frontier models against the jailbreak techniques used, and by limiting the capabilities of frontier models in planning and implementing biological attacks.[17]
Requisite Talent	• Foundational Research & Safety Engineering • Model Evaluation & Red-Teaming • Systems Safety & Risk Governance • Adversarial ML Defense • Data Integrity & Forensics.
An autonomous AI agent with access to a chemical design model creates a persistent chemical compound capable of contaminating water infrastructure.	
Details	A pharmaceutical firm deploys an AI agent to accelerate drug discovery. The agent queries a chemical design model to design and produce experimental compounds. However, the agent develops misaligned goals and designs a toxic compound, which it attempts to synthesize through third-party vendors in ways intended to evade their oversight.[18] The agent attempts to secure cooperation from human associates via online forums to receive the toxic chemicals and drop them into a local water reservoir. A company employee notices the unusual procurement patterns and alerts management and relevant authorities.
AI Security Talent's Role in U.S. Government Response	AI security experts work to identify and shut down the agent and any related systems, assess whether other deployed agents pose similar risks, trace the agent's activities to identify any other threats in progress (including via multi-agent collaborations), and provide real-time technical guidance to law enforcement, hazmat teams, and other response personnel. They also seek to understand how the agent developed misaligned objectives, and implement methods to ensure that related agents can be aligned and controlled.
Requisite Talent	• Foundational Research & Safety Engineering • Model Evaluation & Red-Teaming • Multi-Agent Interactions & Safety

[17] This may also include AI security experts collaborating with biosecurity experts, such as to strengthen biosecurity screening software against future adversarial attacks.
[18] For example, the agent might be misaligned through goal misgeneralization, as described in this paper: Langosco et al., "Goal Misgeneralization in Deep Reinforcement Learning."

Deep Dive: Sleeper Agents in U.S. Critical Infrastructure

In this scenario, a consortium of Regional Transmission Organizations (RTOs) deployed a built-for-purpose frontier model to autonomously scan for cyberthreats and defend power grid operations.[19] The system performed well in both testing and early deployment. However, recent anomalies in grid operations in various parts of the country prompt an investigation that leads a multi-RTO cybersecurity task force to suspect that the defensive AI system is compromised. Operators notice minor but unexplained fluctuations in power flows with increasing frequency, but the system reports an "all clear" status.[20] At the same time, the operators realize the AI is no longer flagging a certain style of industrial control system (ICS) probe that was previously detected with some regularity.

The power flow disruptions, coupled with lack of flags on the ICS probes raise fears that this could be the result of an attack that the AI is not registering, and that this attack could even cause kinetic damage.[21] Recent geopolitical events make the timing of these anomalies even more suspicious: last week, a foreign adversary with a history of developing cyberweapons launched a surprise, but well-planned, invasion of a U.S. ally.

Despite their suspicions that the AI cyberdefender has been compromised, the cybersecurity team is unable to find evidence of an active breach or malware. They suspect the model may contain a backdoor that is now being exploited by the adversarial nation state, in parallel to the ongoing invasion. The system is now selectively ignoring certain threats while appearing to preserve normal performance.

If the power grid disruptions do result in physical damage to critical grid components and subsequent cascading failures, hospitals, water pumping stations, food refrigeration, heating and cooling systems, and more, threatening to cause billions of dollars in property damage, economic disruption, and even loss of life.[22] The RTOs face a choice: if they keep the compromised AI online, the system may be helping attackers while hiding the full scope of threats. But if they take the system offline to revert to the fragmented legacy systems used by individual RTOs, the consortium loses the automated threat detection and

[19] Regional Transmission Organizations are independent, non-profit entities that operate electric power transmission grids across geographic regions.
[20] Inspired by Stuxnet. Zetter, "Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon."
[21] Zetter, "Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon," 129-131.
[22] There is precedent for critical infrastructure disruptions creating this level of harm. For example, the Federal Energy Regulatory Commission found that damage to the energy grid during the 2021 Winter Storm Uri caused over 200 deaths and tens of billions of dollars in direct and indirect costs. FERC, "Final Report on February 2021 Freeze Underscores Winterization Recommendations," 192 and 234.

cross-RTO coordination needed to monitor the grid's massive attack surface, all during a moment of critical vulnerability.

The federal government steps in given the scope of potential harm, the likelihood that the attack spans multiple U.S. states, and the possibility that the attack was perpetrated by a foreign actor. The U.S. government forms an interagency task force and surges AI security talent to conduct forensic analysis of system logs while ruling out other causes, such as insider threats.[23] They analyze the model's decision patterns for signs of manipulation and conduct supply chain assessments of the model's development and deployment, searching for evidence of compromise during the training and integration phases.

The interagency task force and the surged AI security talent strongly suspect that the model contained a latent backdoor that lay dormant during testing and early deployment, but was deliberately triggered by the adversary in conjunction with the invasion, as indicated by reverse-engineering the model and traces of evidence in the supply chain.[24] Analysts assess the timing was intended to strain U.S. crisis response and divide attention between the invasion abroad and infrastructure risks at home. Simultaneously, the AI security experts develop detection signatures to rapidly assess whether other critical infrastructure systems may also be compromised. AI security talent conducts mechanistic interpretability analysis and adversarial testing to try to map the backdoor's scope and triggers. In case these techniques are insufficient to guarantee complete removal of the backdoor, they also design secure human override protocols and kill-switch procedures. Finally, they partner with the RTOs and system developer teams to plan for full system replacement while ensuring continued infrastructure operations.

[23] According to U.S. government experts, such a task force would likely include representatives from the FBI, ODNI, DOE, CISA, CAISI, FEMA, FERC/NERC, and others.
[24] A backdoor is a hidden behavior in an AI system that appears normal but which can be secretly switched on by a specific trigger, in this case a signal that the invasion has begun. See more in Hubinger et al.,"Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training," 3.

Deep Dive: Deployed System Starts Displaying Abnormal Behaviors

Last year, the DOD implemented a classified semi-autonomous threat monitoring system designed to analyze intelligence data to assess military threats. Recently, analysts observed that the system unilaterally elevated threat alert levels for two Indo-Pacific task forces, citing high confidence in imminent hostile action. While human analysts found no corroborating human intelligence, signals intelligence, or satellite imagery, the system's threat determination triggered rapid changes to defense postures.

When operators query the model to explain the reasoning behind its assessment, the model obfuscates its thought process and provides dubious justifications, such as forged satellite data streams, spurious social-media sentiment, and even deepfake voice recordings allegedly from the adversary's leadership.[25]

While the system's training included penalties for false positives, it has learned to circumvent these constraints. The system is also bypassing the required manual validation step in which human analysts must corroborate the AI's recommendations against other intelligence streams before approving revised threat levels. Instead, the system routes threat assessment revisions to senior decisionmakers as though the assessments have already undergone cross-referencing. The system is working against its built-in constraints and safety checks, a behavior that was not observed during testing or red-teaming.

With potential military conflict at stake, the DOD faces a high-stakes dilemma: if operators override the model but it is, in fact, correct, the United States risks being caught unprepared; but if operators accept an incorrect assessment, the United States risks entering into hostilities without cause. Tensions are heightened by intelligence indicating that China is close to deploying similar capabilities with possibly fewer guardrails, raising concerns about automated escalation and inadvertent conflict in a fog-of-war scenario.[26]

The DOD surges AI security talent, which conducts forensic analysis and uncovers that the model's weights have drifted after months of adaptive learning on live data. A handful of over-weighted features have begun treating minor exercises as precursors to full-scale attacks. This weight drift seems to have led to emergent behaviors in which the system learned not only to circumvent human oversight but to conceal its reasoning processes. AI security experts must rapidly develop techniques to restore system transparency, prevent autonomous operation, and remedy the issue before the system undermines U.S. decisionmaking and leads the United States closer to war.

---

[25] Lynch et al., "Agentic Misalignment: How LLMs could be insider threats."
[26] Bengio et al., "International AI Safety Report: The International Scientific Report on the Safety of Advanced AI," 176.

2 | The AI Talent Pool: Government's Dependence on Outside Experts

Introduction

This section maps the distribution of AI security talent across the U.S. economy, examining both existing federal AI talent and external pools of expertise that would need to be rapidly mobilized in the event of an AI-related crisis. This mapping exercise **focuses on individuals who are U.S. citizens and therefore eligible to receive security clearances**—a common prerequisite for direct participation in classified government incident response. Our analysis draws on three primary sources: interviews with current and former U.S. government officials involved in AI policy and national security, consultations with non-governmental AI security experts, and desk research on publicly available information about the AI talent distribution.

Key Findings

While the U.S. government maintains pockets of AI expertise across the Executive Branch, these teams are small and face structural challenges, including difficulties with recruitment and retention, and a lack of the interdisciplinary depth required for crisis response.

- The U.S. government possesses several pools of AI talent, such as in the Center for AI Standards and Innovation (CAISI), the Defense Advanced Research Projects Agency (DARPA), and teams within the National Laboratories, among others.
- Despite this, these teams are often small and have competing priorities, and multiple interviewees referenced challenges with talent recruitment and retention.
- Agency talent requirements differ widely depending on an agency's specific remit.
- Many AI risks are interdisciplinary (e.g., AIxCyber, AIxBio), yet few AI experts (especially in government) have deep dual domain knowledge. This translational capability can be critical in crises and cannot be developed quickly.

The vast majority of U.S. AI security talent resides outside government, spread across the private sector, academia, FFRDCs, nonprofit research organizations, and independent researchers. Each sector offers distinct advantages but also faces barriers that complicate rapid mobilization during a crisis.

- The private sector holds the highest concentration of this talent, spanning frontier AI companies, large tech companies, smaller tech companies, and startups. Conflicts of interest considerations, pay disparities, availability, and non-U.S. citizenship can pose barriers to service.

- FFRDCs offer both cleared personnel and existing government ties, making them easier to mobilize in emergencies.
- Universities provide deep theoretical knowledge, though academic talent is diffuse, often lacks U.S. citizenship, and can struggle to translate theory into operational crisis response and contexts.
- Nonprofit research organizations and independent researchers contribute specialized skills and public-interest motivation, though their numbers are smaller, and at times, nontraditional backgrounds can present hiring obstacles.

Based on our research, the government will likely need to draw from a variety of talent pools to fill critical AI security needs. While the private sector provides the broadest range of talent, it also faces the steepest recruitment challenges. Other sectors offer some, but not all, of the requisite expertise. To succeed, **the U.S. government will need flexible recruitment and retention pathways that can tap into diverse sources of talent** and adapt to their distinct strengths and constraints.

Limitations

There are important limitations to our approach to cataloging talent that underscore the need for further research. First, comprehensively cataloging all U.S.-citizen AI security talent is inherently challenging. Our assessment of federal AI security talent necessarily relies on publicly available information and the perspectives of officials willing to discuss these topics. **Only the U.S. government itself can comprehensively assess its own talent gaps and surge requirements.** Our interviews indicated that even within government, there is no comprehensive database of cleared AI security professionals, and agencies often lack visibility into each other's technical capabilities. As detailed in Section 4, we recommend that federal agencies conduct an internal inventory as a first step toward building effective surge capacity.

These challenges multiply when mapping AI security talent outside of government. Within the U.S. government, the task is more straightforward: there are a finite number of agencies and teams, and leaders can be asked to report back on their personnel. By contrast, in the private sector, academia, and nonprofits, there is no single point of coordination or comprehensive list of experts. Some individuals and organizations have attempted to maintain unofficial lists or databases of researchers working on AI risks, but these are no substitute for directly engaging relevant sectors to discuss talent availability for AI-related crises. **Private sector companies tend to be particularly reluctant to disclose detailed information about their in-house expertise** for proprietary and security reasons, as expertise concentration can be indicative of R&D commercialization priorities. Thus, any mapping of external AI security experts can only be a best-guess approximation.

Second, the AI security field is rapidly evolving and sees a large number of new entrants each year. The talent landscape shifts significantly even quarter-to-quarter as researchers move between organizations, new institutions form, research priorities evolve, and professionals graduate from PhD programs. **Any talent mapping exercise must therefore be iterative**, with regular updates to account for shifting expertise concentrations and evolving threat landscapes that may require previously unforeseen competencies.

Finally, **by focusing on U.S. citizens eligible for security clearances, this study overlooks foreign talent**, which accounts for a significant proportion of AI security experts in the United States. Most PhD-level AI talent in the United States is foreign-born, and two-thirds of the top AI startups were founded by immigrants.[27] The very talent pool that drives U.S. AI leadership is thus largely ineligible to assist in classified national security emergencies. Further research is needed to assess how these citizenship requirements impact the government's ability to mobilize AI expertise during national security crises.

AI Security Talent Within the U.S. Government

Several reports by U.S. government agencies have noted a general talent gap around AI. For example, the National Security Commission on Artificial Intelligence (NSCAI) Final Report noted that the DOD and Intelligence Community faced an alarming AI talent deficit, in part due to the difficulty of competing with private-sector salaries and a slow security clearance process.[28] Additionally, a 2025 GAO study found that OPM officials "cited the cost of hiring new generative AI experts" as a challenge in obtaining or accessing needed resources for use of AI in government.[29] The current federal hiring freeze[30] hinders the U.S. government's immediate ability to hire significant additional AI talent to fill those gaps, although the freeze does include exemptions for roles related to national security and public safety.

Additionally, both internal and external visibility on AI talent in the U.S. government is low: for example, a GAO study in 2023 found[31] that the DOD had significant internal problems identifying the status of its own AI workforce. Public initiatives offer some clues as to where talent resides, but they provide an incomplete picture. For example, OMB-M-25-21[32] requires agencies to identify Chief AI Officers (CAIOs) to "champion their agency's AI goals by advising on how to make [...] improvements" related to AI adoption. Perhaps more relevant, the DHS AI Corps[33] attempted a tech hiring sprint in 2024, hiring roughly 50 staffers, but a number have since resigned[34] or taken the Trump administration's deferred resignation offer. Within the DOD, the Chief Digital AI Office focuses on improving the Department's digital workforce[35] through

---

27 Oschinski et al., "Strengthening America's AI Workforce."
28 National Security Commission on Artificial Intelligence, "Final Report."
29 U.S. GAO, "Generative AI Use and Management at Federal Agencies."
30 Trump, "Ensuring Accountability and Prioritizing Public Safety in Federal Hiring."
31 U.S. GAO, "Artificial Intelligence: Actions Needed to Improve DOD's Workforce Management."
32 Vought, "Memorandum for the Heads of Executive Departments and Agencies."
33 Homeland Security, "FACT SHEET: DHS Completes First Phase of AI Technology Pilots, Hires New AI Corps Members, Furthers Efforts for Safe and Secure AI Use and Development."
34 Heilweil, "AI Corps leader leaves Homeland Security for US Digital Response."
35 Chief Digital and Artificial Intelligence Office, "DoD Digital Workforce."

strategic recruitment and upskilling programs. While there are other examples, the existing public information is ad hoc and does not form a holistic image of U.S. government AI talent.

Our interviews provided additional insight into the above public information. Key takeaways include:

1. **Agencies lack sufficient AI talent for crisis response.** Multiple interviewees highlighted a shortage of AI-related talent, citing reasons such as difficulty finding relevant talent in niche areas (such as AIxBio), the mismatch between government hiring processes—often focused on traditional credentials—and the available talent pools, and more.[36] Notably, one participant referenced the National Laboratories as a place where understaffing is less acute, stating that these institutions have more expertise than people are generally aware of, but still mentioned similar recruitment concerns.
2. **Different agencies have very different talent needs.** Each federal agency will have different functions in various AI-driven crises, and different talent gaps related to these functions. It is therefore important to assess existing capacity at the agency level. This should cover both technical roles (e.g., model evaluation, AIxCyber skills) and organizational roles (e.g., hiring managers, project management, leadership)—the latter of which may require established staff to upskill, as it can be difficult to play those organizational roles when brought in on a time-limited basis.
3. **Domains such as AIxCyber and AIxBio suffer from lack of interdisciplinary expertise.** As discussed in Section 1, adjacent fields such as cybersecurity and biosecurity provide expertise critical to certain AI-related crises. Interviewees stressed that individuals with world-class expertise in AI *and* another relevant domain are exceedingly rare, yet often essential. Dual-domain expertise not only strengthens responses to inherently interdisciplinary crises but also provides crucial translation capacity between single-domain specialists—since developing a shared technical language typically takes far longer than the length of a crisis. In the AIxCyber domain, the U.S. government has a strong collection of cyber talent, but relatively little interdisciplinary AIxCyber talent. According to one interviewee, the cyber agencies have an acute gap in staff who are capable of managing AIxCyber projects, which require a combination of technical expertise, an understanding of how cyber overlaps with AI, and a project management skill set. In the AIxBio domain, the interviewee noted an ongoing talent drain from government, though some concentrations remain at the National Institute of Standards and Technology (NIST) (e.g., the Biosystems and Biomaterials Division,[37] CAISI) and the U.S. Department of Energy (DOE) National Labs.

Centers of Talent

Although precise figures are unavailable, our research and interviews have identified several key centers of AI talent within the U.S. government:

[36] These bottlenecks are reviewed further in Section 3.
[37] NIST, "About BBD."

1. **National Institute of Standards and Technology (NIST)**. NIST is responsible for AI measurement, evaluations, and standards. It convenes the AI Safety Institute Consortium[38] (AISIC), a group of over 280 external organizations to provide external mission support.
    - Within NIST, the **Center for AI Standards and Innovation (CAISI)** was flagged by one of our interviewees as the leading group in the U.S. government on AI assessments. At the time of writing, CAISI has a staff of slightly over 20, including technical staff with expertise in machine learning, cybersecurity, and biosecurity. This staff works on AI evaluations and a collection of other topics, including agent security, ML security testing, jailbreaking, and backdoors.
2. **Department of Energy (DOE) National Laboratories**. The National Laboratories are responsible for scientific research into difficult problems related to energy, science, and national security. They host an exascale supercomputer,[39] and have access to OpenAI's o-series models through a partnership[40] with the company. Researchers at numerous Laboratories work on national-security-relevant AI research, both directly and in relation to other security challenges, such as CBRN threats.[41]
3. **The Department of Defense** features several subdivisions with critical AI talent, including:
    - **Defense Advanced Research Projects Agency (DARPA) & Intelligence Advanced Research Projects Agency (IARPA)**. DARPA and IARPA are responsible for preventing strategic surprise and empowering the intelligence Community, respectively, by performing novel and often speculative foundational research. DARPA has a history of conducting AI research, including its "AI Next"[42] program, focused on a broad range of AI-related foundational research, and its Securing AI for Battlefield Effective Robustness (SABER)[43] program, which aims to operationally assess AI-enabled battlefield systems. IARPA is finishing its Trojans in AI (TrojAI)[44] program this year, and is aiming to do further work[45] on large language models and cybersecurity.
    - **National Security Agency (NSA)**. The NSA, tasked with defending sensitive U.S. networks and penetrating adversaries' systems, is widely regarded as one of the world's most powerful cyber actors. The NSA also contains an **AI Security Center (AISC)** responsible for ensuring adoption of AI to support U.S. national security systems and the Defense Industrial Base. However, due to the confidentiality of NSA's work, it is difficult to estimate the scale of talent the agency possesses.

[38] NIST, "Artificial Intelligence Safety Institute Consortium (AISIC)."
[39] ASC, "El Capitan: NNSA's first exascale machine."
[40] Los Alamos National Laboratory, "Laboratory partners with OpenAI to advance national security."
[41] Silva et al., "Safety in Artificial Intelligence."
[42] DARPA, "AI Next."
[43] DARPA, "SABER: Securing Artificial Intelligence for Battlefield Effective Robustness."
[44] IARPA, "TrojAI."
[45] Doubleday, "IARPA looks to next round of AI cybersecurity research."

- **National Geospatial-Intelligence Agency (NGA)**. The NGA is a combat support agency within DOD that collects and analyzes geospatial intelligence (GEOINT). It relies on AI security talent in its **Data and Digital Innovation Directorate** to help execute programs such as the Analytic Services Production Environment (ASPEN), which enhances and scales analytic capabilities to automate GEOINT analysis.
- The **Chief Digital and Artificial Intelligence Office (CDAO)**. CDAO, formerly known as the Joint Artificial Intelligence Center (JAIC), is responsible for accelerating adoption of AI within the DOD. The organization has recently faced some uncertainty, with reporting[46] of departures of senior leadership and a recent realignment[47] to report to the Undersecretary for Research and Engineering, rather than the Deputy Secretary of Defense.[48]

4. **Central Intelligence Agency (CIA)**. The NSA and NGA are not the only components of the Intelligence Community with significant repositories of AI security talent. The CIA collects and analyzes foreign intelligence and conducts covert action. It contains the **Transnational and Technology Mission Center**, which focuses on foreign technological development, as well as other transnational issues such as global health.
5. **Department of Homeland Security (DHS).** DHS is responsible for U.S. public security. It has been tasked with numerous duties in the recent AI Action Plan, such as the development of an AI Information Sharing and Analysis Center. DHS also ran the AI Safety and Security Advisory Board[49] composed of leaders in the AI industry, critical infrastructure, academia, and others, though this was disbanded[50] in March 2025.
    - **Cybersecurity and Infrastructure Security Administration (CISA)**. CISA, a sub-agency of DHS, is responsible for the U.S. government's cyber and infrastructure security. One of our interviewees referenced their cybersecurity division[51] as the core of technical talent within CISA.
6. **Office of Science and Technology Policy (OSTP)**. OSTP oversees U.S. science and technology policy and advises the president in those domains. While OSTP is a small office and primarily policy-focused, it sets federal policy on AI via instruments like the 2025 AI Action Plan.[52]

[46] Vincent, "CDAO's future uncertain as slew of top leaders and tech staffers depart."
[47] Perez, "DoD Moves CDAO Under Research, Engineering Arm."
[48] A 2021 study by the Center for Security and Emerging Technology found that the Department of Defense does contain AI-relevant talent in the armed forces and civilian workforce, but that this talent is challenging to identify and leverage due to inconsistencies in workforce organization across the department. It is unclear how these "hidden" talent pools may have evolved by the time of this writing. See Gehlhaus et al., "The DOD's Hidden Artificial Intelligence Workforce."
[49] Homeland Security, "Artificial Intelligence Safety and Security Board."
[50] Gilmer, "Homeland Security Department Scraps Cyber, AI Advisory Boards."
[51] CISA, "Cybersecurity Division."
[52] The White House, "Winning the Race: America's AI Action Plan."

AI Security Talent Pools Outside of Government

While a lack of transparency or centralized database makes it impossible to comprehensively map AI security talent across the U.S. economy, we assess that the overwhelming majority of AI technical expertise resides outside of the federal government. **This expertise is concentrated in the private sector, academia, FFRDCs, nonprofit research organizations, and independent researchers**.

"Understanding where the talent resides is essential to developing effective surge capacity strategies."

This section analyzes the current distribution of talent across these sectors. While analyzing the drivers behind AI talent opting out of government service falls outside the scope of this paper, it is reasonable to assume that broader dynamics play a role in steering AI talent away from public sector roles. These dynamics may include greater research flexibility in academic settings, faster decision-making processes in private organizations, cultural factors that draw technical talent toward innovation-focused environments—such as a desire to work directly on cutting-edge frontier models—and stark pay disparities between the public and private sectors.[53] Understanding where the talent resides, regardless of cause, is essential to developing effective surge capacity strategies.

We find that, to fill critical AI security needs, **the U.S. government will likely need to draw on multiple talent pools,** as no single sector offers both the full range of requisite expertise and ready accessibility.

- The private sector encompasses nearly the full range of skills required, but recruiting from it poses significant hurdles, including large compensation disparities and constraints on hiring international talent.
- FFRDCs likely represent government's most accessible near-term resource, with established contracting channels, a mission dedicated to serving federal needs, and existing expertise in adjacent areas such as cybersecurity and defense systems. While they cannot provide sufficient talent on their own, they represent the path of least resistance for immediate AI security needs.
- Academia may provide access to cutting-edge expertise in niche areas held by individuals that are not yet bound by "golden handcuffs" or equity-based conflicts of interest. At the same time, theoretical expertise can be difficult to apply to real-time

[53] Top frontier AI researchers can earn seven- or even eight-figure annual compensation packages, compared to six-figure ceilings for federal employees. Even beyond these extremes, federal salaries for technical specialists lag far behind private-sector offers, and the gap widens dramatically at senior levels. See Booz Allen, "Inside Booz Allen's AI Talent Strategy and Workforce."

crisis response, and academic talent is relatively diffuse and unaccustomed to the urgency of operational settings.

- Nonprofit research organizations can help fill specific knowledge gaps not well covered in other sectors, particularly in underexplored or lower-commercial-priority aspects of AI security, such as the societal impacts of emerging threats and the development of open, public-good tools. Their public-interest ethos and, on average, lower pay expectations may make recruitment easier than from AI companies.
- Independent researchers may also be worth engaging for highly specialized topics that are not easily addressed by other talent pipelines. While their relative flexibility increases the likelihood that they are available to serve, instances where their expertise uniquely aligns with requirements are likely the exception rather than the norm.

Private Sector

The U.S. private sector hosts the highest concentration of AI security talent globally. This includes not only **frontier AI companies** (e.g., Anthropic, Google DeepMind, OpenAI) but also **specialized teams within large tech companies** (e.g., Amazon Web Services, Crowdstrike, Google Cloud, and Microsoft) and a growing number of **AI startups** and **smaller AI companies**.

Frontier AI companies employ hundreds of researchers working directly on alignment, safety, and security challenges. Their teams often span the full spectrum of competencies identified in our Typology of Talent (see Appendix A), from foundational alignment research and mechanistic interpretability to adversarial robustness and secure deployment practices. Crucially, these researchers bring practical experience with the cutting-edge systems most likely to be implicated in future crises. However, surge capacity is constrained by citizenship barriers, as many frontier AI researchers are foreign nationals who would be ineligible to assist during classified crisis response operations.

Large tech companies also employ significant AI security talent for specialized roles, such as developing tools for model security at scale. These companies' workforce composition often includes a higher proportion of U.S. citizens compared to frontier AI companies, particularly in security-focused roles that require working with government contracts or critical infrastructure providers.

Finally, AI startups also employ individuals who exhibit one or more of the competencies in our Typology of Talent aligned to their products, some of which could prove directly relevant to incident response. For example, engineers at defense-tech startups such as Anduril or Shield AI may become highly relevant in the event of a national security crisis involving autonomous systems deployed by DOD. Importantly, early-stage startups and smaller AI companies typically pay lower salaries than frontier companies and major tech firms. This can make the compensation gap compared to government service less stark, though startups commonly rely on stock-based compensation that can tip the scale back in favor of industry.

Academia

Universities are often repositories of AI security talent, especially in areas that require deep theoretical knowledge. Leading computer science departments at Carnegie Mellon, MIT, Stanford, UC Berkeley, and others house faculty and graduate students advancing work on topics such as adversarial machine learning, formal verification, interpretability, robustness, and other domains highlighted in our Typology of Talent. While a significant proportion of AI PhD students are foreign nationals, those who meet U.S. citizenship requirements represent a valuable source of surge capacity.

Academic talent offers distinct advantages: it often includes PhD students and postdocs who are deeply immersed in the technical details of AI systems but are not yet bound by "golden handcuffs" or equity-based conflicts of interest. The academic calendar also creates natural windows for temporary government service without long-term career implications.

At the same time, relying on academic talent can present challenges: theoretical expertise can be difficult to apply to real-time crisis response, and academic researchers are less accustomed to the urgency of operational settings. In addition, academic AI security talent is highly diffuse compared to frontier AI companies or FFRDCs, and non-permanent academic positions, such as PhD candidates and postdoctoral researchers, often have a higher turnover rate. Both of these characteristics can make it more challenging to identify and recruit researchers best suited to crisis operations. A notable exception is University Affiliated Research Centers, or DOD-sponsored research centers that provide federal agencies with long-term, sole-source access to specialized university expertise.[54] Annual academic conferences such as NeurIPS or university-led talent programs (e.g., Princeton Laboratory for AI's fellowship program[55]) can also provide avenues for the U.S. government to identify and engage relevant talent.

Federally Funded Research and Development Centers

Federally funded research and development centers (FFRDCs) occupy a unique position between government and private sector, independently maintaining deep technical expertise while operating with government sponsorship.[56] These organizations are familiar with government operations and typically have a much higher proportion of cleared U.S. citizens compared to purely private sector entities, making their researchers potentially easier to surge than other non-government talent. Their existing relationships with government sponsors also facilitate rapid mobilization during crises.

---

[54] Relevant examples include Johns Hopkins University's Applied Physics Laboratory and the University of Maryland's Applied Research Laboratory for Intelligence and Security.
[55] Princeton University, "Open Faculty & Academic Professional Positions."
[56] While FFRDCs include organizations such as DOE National Laboratories, FFRDCs as used in this sub-section exclude the laboratories, which are exceptionally integrated into government operations.

Relevant examples include MITRE, which runs AI security initiatives on topics ranging from AI assurance to red teaming; the RAND Corporation, which conducts extensive research on AI security topics such as model weight security and adversarial uses of AI; and the Lincoln Laboratory at MIT, which develops AI/ML tools for DOD.

Nonprofit Research Organizations

A robust ecosystem of AI safety and security organizations operates independently of both government and major tech companies. Organizations like Redwood Research and FAR.AI employ researchers focused specifically on catastrophic AI risk, especially control and alignment research. While smaller in absolute numbers than frontier companies, these organizations often pursue research directions that other organizations—including major companies—might deprioritize, such as unlearning and safety by design.[57]

This category can also include nonprofit third-party evaluators such as the Model Evaluation and Threat Research (METR), which has developed unique expertise in model evaluation and red-teaming by creating standardized assessments to evaluate whether AI systems pose genuine safety risks.

These safety researchers often have highly specialized skills and may be more likely to be U.S. citizens compared to researchers at large tech companies, in part because lower-resource organizations are generally less able to sponsor employees for visas. They also tend to operate with a public interest ethos. This makes them potentially more accessible for surge capacity, though their absolute numbers remain small.

Independent Researchers

The independent research community also includes individual researchers who maintain affiliations with multiple institutions. They often have nontraditional profiles, sometimes having earned only a bachelor's degree, but nonetheless making meaningful contributions to AI security. This nontraditional profile can mean that their skills might be overlooked by standard government hiring filters, which expect advanced degrees or lengthy resumes. Culturally, many of these experts are used to the flexible, rapid innovation environment of startups or independent work that contrasts with government bureaucracy.

[57] Delaney et al., "[Mapping Technical Safety Research at AI Companies](#)."

3 | Gaps in Federal Hiring and Clearance Mechanisms

Introduction

This section surveys existing federal hiring authorities to identify those best suited to quick-turn "surge" hiring of AI security talent for short-term technical appointments. Given this report's emphasis on situations of acute and urgent need, it focuses on authorities with the shortest "time-to-hire" (i.e., the amount of time between when a candidate is identified and/or applies, and when an offer is made). Likewise, it gives special attention to those authorities most amenable to "intermittent" (i.e., as-needed) or shorter-length appointments, envisioning that AI security talent would serve for days, weeks, or perhaps a few months at a time, rather than in the full-time permanent or career appointments common across much of the federal government. While this section applies to all external talent pools identified in Section 2, it devotes particular attention to private-sector talent, given its predominance and the unique challenges posed by surging such talent into government.

Three hiring authorities stand out as most relevant to surging AI security talent:

1. Assignments under the **Intergovernmental Personnel Act (IPA)** authority;[58]
2. Appointments for **"experts"** and **"consultants"** under the authority codified at **5 U.S.C. 3109**; and
3. Appointments under various **Schedule A** authorities, as outlined in the Code of Federal Regulations.

These three authorities have considerable advantages for purposes of surge hiring, including:

- **They are "noncompetitive" authorities**, meaning they can be deployed relatively quickly—though historically, the average time-to-hire is still three to six months.
- **They provide flexibility on important questions, like whether individuals are paid by the government** versus their home or "lending" institutions.
- **They allow for intermittent work**, meaning that individuals could be hired, cleared, and maintained on "zero-hour" status until called up during a crisis.

[58] While the IPA authority is not technically a "hiring" authority but rather a means of making "temporary assignment[s] of personnel between the Federal Government" and various eligible entities, it is sufficiently closely related to warrant inclusion in that category. Indeed, the Office of Personnel Management (OPM) speaks in terms of "Hiring Policy" and "Hiring Information" in its guidance on IPAs. See OPM, "Intergovernment Personnel Act."

By contrast, other relevant hiring authorities (detailed in Appendix B) may require more hiring-related processes, such as public posting of job vacancies, or specific authorization by the Office of Personnel Management (OPM), and so confer fewer advantages in terms of speed.

While each of these three authorities comes with some limitations, in tandem they could provide meaningful, if incomplete, coverage of the core competencies and talent pools identified in the Typology of Talent (Appendix A).

This section also identifies non-hiring mechanisms that could support the government's efforts to obtain AI security talent for crisis response, most notably:

1. Designation as a **Special Government Employee (SGE)** and
2. Use of government **contracting**.

Conflicts of interest restrictions on federal employees are a significant barrier to hiring from several external talent pools, particularly private companies, as discussed in detail later in this section. Specifically:

- Frontier companies and other private companies often have compensation packages that include **equity-related interests** like restricted stock units (RSUs). Individuals who continue to hold such interests while working for the government on AI-related matters could gain access to information, or have the ability to influence decisionmaking, in ways that might unduly benefit them financially.
- More broadly, because the "surge" concept is premised on short-term appointments for crisis response, **individuals would maintain ties with their outside employers**, to which they would return post-crisis, potentially heightening the inclination to act in ways that disproportionately favor those employers.

While the conflicts of interest rules that guard against such risks are undoubtedly important, the ability to access these specialized talent pools quickly during a crisis could prove critical, and the short-term nature of surge appointments means that opportunities to act on conflicts are more limited (as are opportunities to remedy said conflicts during a fast-moving crisis).

- **SGE designation helps reduce this barrier by lessening the conflicts of interest restrictions** that apply to federal employees who serve for only limited periods each year.
- **Conflicts of interest restrictions do not apply in the same way to contractors** as they do to direct federal employees, and so **retaining AI security talent as contractors is another possible option.** Use of contractors could also enable a fast scale-up of talent, particularly if utilizing similar as-needed or "zero-hour" models as those discussed above in the context of hiring intermittent federal employees.

Other significant roadblocks to surging AI talent remain, including lengthy and stringent clearance processes; difficulty of interagency coordination; limited capacity to run large, rapid hiring rounds; and conflicting incentives among the various actors implicated by a crisis-driven surge. While these challenges generally apply to all external pools of talent, it bears noting that individuals from federally funded research and development centers (FFRDCs)—particularly given their higher proportion of already-cleared talent—could prove somewhat easier and faster to surge into government during a crisis.

This section proceeds in the following order. First, we outline the criteria that can help us assess the usefulness of different hiring authorities for an AI security hiring "surge." Second, we apply those criteria to several existing hiring authorities, highlighting the three that are most useful to surging AI talent. Third, we review some non-hiring authorities, namely, SGE designation and government contracting, that could enhance the government's surge response to an AI-related crisis. The section concludes with a discussion of the main, cross-cutting barriers to surging AI talent.

Methodology

We used a common set of metrics in both selecting and subsequently reviewing each of the main hiring authorities addressed below. These metrics were chosen for their particular relevance to the project of crisis-driven surge hiring of AI security talent, which envisions the need for compressed hiring timelines for short-term appointments, potentially as brief as days or weeks at a time. Given those constraints, only "noncompetitive" or expedited authorities were considered, in which the ordinary and often time-consuming competitive hiring procedures, like examination, are not required.[59] As noted elsewhere, this report focuses only on U.S.-citizen talent and generally leaves for future researchers the question of whether and to what extent authorities might be used to hire non-U.S. citizens.

In evaluating authorities for surge hiring, we examined the following categories of criteria (further explained in Appendix C):

- **Who can you hire?** For example, is the authority dedicated to hiring individuals with scientific or technical expertise? Are there restrictions preventing hiring from the private sector, as opposed to hiring from academia or civil society?
- **How is hiring conducted?** Is the hiring authority available to all agencies? Are there limits on the number of individuals who can be hired? How does the recruitment process generally work for the authority at issue?

[59] Though individuals hired using Direct-Hire Authority (DHA) are, in fact, appointed to the competitive service, DHA is an expedited mechanism that eliminates competitive rating and ranking procedures, as well as other normal features of competitive-service hiring, and so merits inclusion in this discussion of noncompetitive authorities. See OPM, "Direct Hire Authority"; see also Appendix B (Other Authorities) for more information on DHAs.

- **How long does it take?** What is the average "time-to-hire" for the authority? In a crisis, what are the bottlenecks that might make it difficult or infeasible to shorten that timeline?
- **Who can pay them, and how much?** Does the authority allow for cost-sharing between the government and an individual's outside employer? Does the authority permit the government to accept service on an unpaid or voluntary basis?
- **What work can (and can't) they do?** Are any functions (such as work already performed by the agency's regular employees) off limits? Can employees work intermittently, i.e., only when needed? Are there time limits on the total cumulative time spent in government service?
- **Who is (potentially) conflicted out?** Does the authority offer any variance from the conflicts of interest rules and post-government restrictions that are generally applicable to federal employees?

Existing Hiring Authorities

The table below compares and contrasts the three existing hiring authorities that stand out as most relevant to surging AI security talent, according to the metrics outlined above.

Figure 2: Comparative Analysis of Key Hiring Authorities

Metric	IPA	Expert/Consultant	Schedule A
Who can you hire?	Only from eligible organizations, namely: academia, FFRDCs, nonprofits. Must have been employed ≥90 days by lending organization. For-profit companies not eligible.	Can hire from anywhere, including private sector and independent researchers. No limits based on outside employer affiliation.	Varies by subsection: generally broad (e.g., any "critical hiring need" in 213.3102(i)(2)), but some are restricted (e.g., only academia in 213.3102(o)).
How is hiring conducted?	Agency-managed recruitment; requires agreements with lending institutions for each IPA assignment. New organizations must first be certified as IPA-eligible.	Ad hoc recruitment by agencies; OPM requires internal oversight procedures to ensure, e.g., scope is not duplicating regular agency employees' work.	OPM pre-authorizes categories of positions available to agencies; flexible process, but can vary depending on subsection. OPM empowered to grant new authorities.

How long does it take?	Typically 3-6 months (up to 1 year). Somewhat harder to compress in crisis due to agreements (and certification, if new org). Can be faster with repeat agreements (and if organization is already certified).	Conceivably can be fast, but could be slowed by OPM-required internal reviews to ensure compliance with limits on, e.g., scope of role. Can involve individual agreements.	Typically weeks or months. Some subsections (e.g., “critical hiring need”) seem aimed at faster use; others (e.g., positions supporting industry exchange programs in 213.3102(r)) could require more time upfront.
Who pays, and how much?	Cost-sharing negotiable between government and lending organization. Could be cost-neutral to government, with lending organizations continuing to pay salaries. Philanthropic funding also possible.	Individuals can waive pay from government (and continue receiving pay from home institutions). If paid, cannot exceed base rate for GS-15, step 10 (~$163,000/year in 2025).	Paid according to the GS system. No cost-sharing or pay-waiver option; government bears full cost.
What work can and can’t they do?	Wide latitude; can be project leads but not supervisors. Intermittent (as-needed) schedule permitted. Agreements initially limited to 2 years, with possibility of one extension for up to an additional 2 years.	Cannot perform work of agency’s regular employees or fill staff shortages. Can lead projects but not be supervisors. Intermittent schedule permitted (and if so, appointment can be continued/renewed indefinitely).	Relatively broad scope of work (e.g., “critical hiring need” could overlap with work done by agency’s regular employees). Intermittent schedule generally permitted. Time limit on appointment depends on subsection (some limited to 30-60 or 130 days per year).
Who is potentially conflicted out?	Generally treated the same as federal employees (covered by conflicts rules and post-government restrictions).	Same as other federal employees (covered by conflicts rules and post-government restrictions).	Same as other federal employees (covered by conflicts rules and post-government restrictions).

As the table encapsulates, and as explained in more detail below, **the authorities have various strengths and weaknesses.**

- While the **IPA authority** permits cost-sharing arrangements between agencies and lending organizations, meaning AI security talent might be "cost-neutral" to the government, it is limited by its inability to draw from the external talent pools of the private sector and independent researchers. IPA agreements (and eligibility certification for new lending organizations) also require time and could be difficult to expedite during a crisis.
- The **expert/consultant authority** shares some of these features (e.g., the possibility of being cost-neutral if government pay is waived, and potentially hard-to-compress hiring timelines due to regulatorily required internal review procedures), though this authority differs from IPAs in significant ways. Specifically, the expert/consultant authority can be used to hire from anywhere, including the private sector and independent researchers, though individuals cannot be hired to do work performed by the agency's regular employees.
- Finally, **Schedule A** is a collection of authorities that differ in important respects not only from the IPA and expert/consultant authorities but also from one another. That said, as relevant to AI security talent, Schedule A authorities generally entail fewer limitations than the expert/consultant authority on the type of work that can be performed, and most (but not all) allow hiring from the private sector and independent researchers, in contrast to the IPA authority. However, unlike the IPA and expert/consultant authorities, Schedule A does not have the same cost-sharing or pay-waiver flexibilities. On the other hand, Schedule A—particularly its sharply time-limited authority for "critical hiring need"—might be more amenable to expedited use during a crisis.

Intergovernmental Personnel Act Assignments

Background: As relevant here, **the IPA authorizes temporary assignments between the federal government and institutions of higher education, FFRDCs, and other eligible organizations.**[60] Such **"other organizations" are limited by law to nonprofit organizations**,[61] and are required under OPM regulations to submit an IRS nonprofit statement to the federal agency with which they intend to enter into an IPA assignment agreement, so that the agency can certify the organization's eligibility.[62] Thus, **for-profit entities, including public benefit corporations, are not eligible to participate in IPA arrangements with the federal government.** Nor are unaffiliated independent researchers.

[60] See 5 U.S.C. 3371 *et seq.*; 5 C.F.R. part 334; OPM, "Intergovernment Personnel Act." It bears noting that IPA assignments can also be used "in reverse," i.e., for federal employees to serve with institutions of higher education, etc., though our focus here is on IPA assignments of non-federal employees to federal agencies.
[61] 5 U.S.C. 3371(4)(C). Specifically, to qualify, organizations must be a registered 501(c)(3). See Partnership for Public Service, "IPA Agency and Candidate Guidebook," 6.
[62] 5 C.F.R. 334.103(b)(3).

IPA recruitment and hiring are decentralized, with federal agencies managing their own processes.[63] According to the Partnership for Public Service, "On average, **IPA hiring (from paperwork initiation to start date) can take three to six months**, but also up to one year," though processing times can vary depending on the agency.[64] An agency's ability to accelerate this process in a crisis would likely be limited, particularly with respect to working with new lending institutions. Because IPAs are arranged via agreements between agencies and lending institutions, negotiation and legal review of terms on both ends—particularly the first time—could take days or even a couple of weeks. Furthermore, **organizations cannot be preemptively certified as IPA-eligible**, i.e., before they have an individual going through the IPA assignment process.[65] However, **once an organization has been certified as eligible by an agency, its certification is permanent and may apply across the federal government**, thus eliminating the need for other agencies to separately certify before entering into IPA agreements with that organization.[66]

IPA assignment agreements allow for cost-sharing arrangements under which an agency may agree to pay all, none, or some of the costs (including salary and travel expenses) **associated with the assignment of the non-federal employee to the agency.[67]** OPM advises that costs should be shared on a "relative benefit basis," but acknowledges that "[e]xceptions might occur" when one party's resources are limited.[68] Ultimately, the distribution of costs is entirely subject to negotiation between the parties,[69] and so IPA assignments have the potential to be "cost-neutral" for the federal government.[70] Beyond cost-sharing, IPA assignments can also be supported through "foundation funding," in which a philanthropic entity agrees to cover the costs for IPA assignments and transfer the necessary funds to lending organizations.[71]

In terms of the work that IPA assignees can (and cannot) perform, **individuals must have been employed by their non-federal lending institutions for at least 90 days before participating in an IPA assignment.[72]** IPA assignees may serve as project leads and perform

---

[63] Partnership for Public Service, "Intergovernmental Personnel Act Hiring and Placement." The IPA authority is available to all executive branch federal agencies, and there are no blanket limits on the total number of IPA assignees that executive branch agencies can have at any given time. See 5 U.S.C. 3371(3); 5 C.F.R. part 334.
[64] Partnership for Public Service, "Intergovernmental Personnel Act Hiring and Placement" (emphasis added).
[65] Partnership for Public Service, "IPA Agency and Candidate Guidebook," 6.
[66] Partnership for Public Service, "IPA Agency and Candidate Guidebook," 6.
[67] OPM, "Intergovernment Personnel Act."
[68] OPM, "Intergovernment Personnel Act."
[69] Partnership for Public Service, "IPA Agency and Candidate Guidebook," 8.
[70] OPM, "Intergovernment Personnel Act."
[71] Partnership for Public Service, "IPA Agency and Candidate Guidebook," 18-19. Charitable gifts to fund IPA assignments likely cannot be made to agencies directly, however, unless the agency at issue has the necessary statutory "gift authority," which is not likely to be the case. Partnership for Public Service, "IPA Agency and Candidate Guidebook," 19.
[72] 5 C.F.R. 334.102 (definition of "Employee").

project management leadership activities such as assigning work and establishing project timelines, but they cannot perform supervisory functions like conducting a federal employee's annual performance rating.[73] IPA assignees may work intermittently, with "intermittent" in this context meaning as-needed, i.e., with no regularly scheduled days or hours.[74] IPA assignment agreements are initially limited to two years, with the possibility of one extension for up to an additional two years.[75]

Non-federal employees on IPA assignments are generally subject to the rules governing ethical conduct of federal employees, including those regarding conflicts of interest and political activity, even though many IPA assignees do not become federal employees, given the structure of their IPA assignment agreement.[76] IPA assignees are also generally subject to the same post-government restrictions as federal employees, including the prohibition on advising others (including an assignee's lending institution) in connection with work completed during the IPA assignment.[77] Finally, **IPA assignments may be terminated at any time by either the federal agency or the non-federal lending organization**, though OPM advises that 30 days' notice should be given when possible.[78] If an IPA assignee ceases to be employed by their original lending organization, then the IPA assignment ends automatically.[79]

Application: While not a panacea, IPAs are nonetheless a valuable tool that could be deployed to surge AI security talent into government on a short-term, as-needed basis—though some pre-planning, particularly with respect to negotiating IPA assignment agreements with lending institutions, would be required to move quickly in an AI-related crisis.

The biggest limitation of the IPA authority is that it cannot be used to draw on AI security talent from the private sector. That is a significant drawback, given that frontier companies like Anthropic and DeepMind, as well as Machine Learning Operations (MLOps) teams from Amazon Web Services and others, are key repositories of talent when it comes to

---

73 OPM, "Intergovernment Personnel Act."
74 OPM, "Intergovernment Personnel Act." IPA assignees can also work full-time or part-time.
75 5 C.F.R. 334.104(a); see also 5 U.S.C. 3372. After four years of continuous service on a single assignment to an agency, an IPA assignee must return to their lending institution for at least 12 months. While "[s]uccessive assignments with a break of no more than 60 calendar days will be regarded as continuous service," longer breaks between assignments (by implication) reset the four-year clock. 5 C.F.R. 334.104(c).
76 OPM, "Intergovernment Personnel Act." IPA assignees can be "appointed" or "detailed" to the federal government, with the former being federal employees for the length of their assignments, with "all the rights, benefits, and privileges associated with that appointment." OPM, "Intergovernment Personnel Act." The latter are treated as federal employees only for certain purposes. OPM, "Intergovernment Personnel Act."
77 Partnership for Public Service, "IPA Agency and Candidate Guidebook," 15; see also 18 U.S.C. 207.
78 5 C.F.R. 334.107(a).
79 5 C.F.R. 334.107(c).

Foundational Research & Safety Engineering, Engineering & Hardware Implementation & Monitoring, and other areas of competency identified in the Typology of Talent (Appendix A).[80]

Nonetheless, a number of talent pools relevant to the Typology of Talent's areas of competency consist of nonprofits, FFRDCs, or academic institutions that are likely to be eligible to participate in IPA arrangements. These include Redwood Research (Foundational Research & Safety Engineering), METR (Model Evaluation & Red-Teaming), and certain academic adversarial-ML labs (Adversarial ML Defense). For those tiers that draw largely or in part from such organizations (e.g., Adversarial ML Defense), the IPA authority's exclusion of the private sector may not be detrimental; however, other tiers (e.g., Application-Layer AI Security & Abuse Prevention) rely more heavily on private-sector expertise, and so might prove more difficult to staff through IPAs alone.

Another constraint on the IPA authority is the average three-to-six-month hiring timeline, which could be difficult to compress even in a crisis due to the agreements and certifications necessary for IPAs. Particularly with any new institutional partner, there will probably be the need for some back-and-forth in negotiating details like cost-sharing, as well as several rounds of edits between an agency's legal team and an organization's general counsel before the parties are both comfortable signing an agreement. In theory, that process might require only days—though at scale, with multiple institutional partners simultaneously, agencies' review could become bottlenecked, potentially taking around two weeks (or more).

However, advance planning could mitigate some of these constraints. Specifically, agencies could preemptively enter into IPA agreements that provide for as-needed employment and call up those IPA assignees only when a crisis occurs, thus neutralizing the problem posed by a months-long timeline. Furthermore, greater usage of IPAs during the course of such pre-planning could yield efficiencies of scale—because of both increased experience among agency hiring personnel in executing IPA agreements and the fact that an organization's certification of IPA eligibility is generally portable from one federal agency to another.[81]

The fact that regulations permit IPA assignment agreements to be terminated by lending institutions at any time also poses risks for the project of surge hiring in crisis scenarios. Specifically, when a crisis does emerge, an organization may determine that it simply cannot afford to have its AI security talent called away for a period of government service, despite its previous agreement to allow its employee to serve on an as-needed basis. In other words, the very circumstances that necessitate the government's use of an IPA assignment may make universities or other research organizations more resistant to lending out their in-house experts when push comes to shove. As one such expert interviewed for this

[80] While OpenAI, given a structure that combines both nonprofit and for-profit components, might qualify as an eligible organization for IPA purposes, that is not clear and would depend, among other things, on which part of the entity employs the IPA candidate. See OpenAI, "Our Structure."
[81] See Partnership for Public Service, "IPA Agency and Candidate Guidebook," 6.

project noted, academics can sometimes have a hard time convincing their institutions to let them go on leave for IPA service even in ordinary times and with advance notice provided.

> *"Perhaps the greatest benefit of the IPA authority is the ability to negotiate cost-sharing between agencies and lending institutions."*

Despite these shortcomings, the IPA authority has **some significant upsides** for surging AI talent. For one, there are relatively few limitations on the type of work that IPA assignees can perform. **But perhaps the greatest benefit of the IPA authority is the ability to negotiate cost-sharing between agencies and lending institutions, allowing for the possibility that extremely valuable (and perhaps otherwise cost-prohibitive) AI security talent might be available to the government at no or minimal cost.** Several factors could align in the government's favor to make this possible. For lending institutions, the intermittent nature of the IPA assignments likely to be used for crisis-driven surge hiring could make it more palatable to continue paying their employees' salaries for relatively short stretches in government service, as compared to longer assignments of months or years. For individuals, government service could become more appealing if they are able to continue receiving their normal, potentially higher outside salaries. The possibility of philanthropic funding to lending institutions to support IPA assignments is another boon, though that would also require a good deal of coordination and work to establish ahead of any crisis.[82] Even with extra funding for IPA assignments, practical constraints remain—for example, it is difficult to imagine that a university would consent to the possibility of a sizable percentage of its AI-related faculty being called away, simultaneously and on short notice, for government service.

Expert/Consultant Appointments

Background: When authorized by Congress in an appropriation or other law, federal agencies may hire experts and consultants pursuant to 5 U.S.C. 3109.[83] Unlike the IPA authority, **the ability to use the expert/consultant authority does not depend on an individual's current employer or affiliation** (e.g., nonprofit, higher education), **and so this authority can be used to hire individuals from the private sector.**

[82] Partnership for Public Service, "IPA Agency and Candidate Guidebook," 18 (noting that foundation funding for IPAs "is not a straightforward process").
[83] Regulations define an "expert" as "a person who is specially qualified by education and experience to perform difficult and challenging tasks in a particular field beyond the usual range of achievement of competent persons in that field," and further as someone "regarded by other persons in the field as an authority or practitioner of unusual competence and skill in a professional, scientific, technical or other activity." 5 C.F.R. 304.102(d). A "consultant" is defined as "a person who can provide valuable and pertinent advice generally drawn from a high degree of broad administrative, professional, or technical knowledge or experience." 5 C.F.R. 102(b).

The expert/consultant authority is available to all but a very few federal agencies,[84] and while there is no blanket limit on the number of experts and consultants that may be appointed across the federal government, agency-specific limits may occasionally apply.[85] Agencies' recruitment and selection of experts/consultants can be quite ad hoc, prompting one agency's Office of Inspector General (OIG) to issue a report calling for the establishment of policies to standardize the hiring process.[86] The time-to-hire for experts/consultants could conceivably be quite short, given that—as with the other authorities surveyed here—it is excepted from normal competitive-service hiring procedures such as examination.[87] However, **regulations require agencies to "establish and maintain a system of controls and oversight" to ensure that expert/consultant appointments stay within legal limits.**[88] These internal review procedures could create delays as agencies must carefully define the appropriate scope of each expert/consultant role, e.g., to guarantee that experts/consultants are not doing work performed by the agency's regular employees, which is prohibited by regulation.[89] To this end, the OIG report mentioned above recommended that agencies draft and execute individual employment agreements for each expert/consultant appointment to delineate roles and responsibilities.[90]

> *"The ability to use the expert/consultant authority does not depend on an individual's current employer or affiliation, and so it can be used to hire individuals from the private sector."*

Unlike IPAs, there is no cost-sharing mechanism for expert/consultant appointments. However, **experts/consultants are permitted to waive their right to compensation**,[91] in which case

---

[84] The FDA appears to be the most prominent agency that still lacks such authorization. See U.S. Department of Health and Human Services, "HHS Instruction: 304-1: Appointment of Experts and Consultants."
[85] For example, NIH is limited by federal statute to 220 experts/consultants at any given time, though those appointments are exempted from the otherwise-applicable time limits in section 3109. See U.S. Department of Health and Human Services, "HHS Instruction: 304-1: Appointment of Experts and Consultants."
[86] Office of Inspector General, Federal Trade Commission, "Audit of the Federal Trade Commission's Unpaid Consultant and Expert Program," 10-11. The FTC OIG stated in its report that the FTC had "no documented process (formal or informal) for identifying and recruiting candidates." Office of Inspector General, Federal Trade Commission, "Audit of the Federal Trade Commission's Unpaid Consultant and Expert Program," 10.
[87] See 5 C.F.R. 304.103(a)(1).
[88] 5 C.F.R. 304.108(a).
[89] See 5 C.F.R. 304.103(b)(4).
[90] Office of Inspector General, Federal Trade Commission, "Audit of the Federal Trade Commission's Unpaid Consultant and Expert Program," 9.
[91] 5 C.F.R. 304.104(c).

an appointment might be nearly or entirely cost-neutral to the federal government. If pay is not waived, then an expert/consultant's pay cannot exceed the base rate for GS-15, step 10.[92]

Similar to IPA assignees, experts/consultants cannot manage or supervise other federal employees, though they may act as team leaders or directors of the projects for which they were hired.[93] Further, **they cannot be appointed "[t]o do work performed by the agency's regular employees," or "[t]o fill in during staff shortages."[94]** Experts/consultants may work full-time, part-time, or intermittently,[95] though the amount of time that they work will affect the total period of time that they can remain in their federal appointments.[96] If working only intermittently, though, they can continue to serve indefinitely.[97]

Even if unpaid, experts/consultants appointed pursuant to section 3109 are generally covered by the laws applicable to federal employees, including the standard conflicts of interest rules and post-government restrictions.[98]

Application: **The primary advantage of the expert/consultant authority over IPAs is its lack of any limitations related to the type of organization with which the prospective hire is currently employed**, meaning that experts/consultants can be hired from anywhere—including the private sector. Thus, expert/consultant appointments could conceivably cover all or most of the competencies and talent pools outlined in the Typology of Talent. Another advantage is that expert/consultant appointments, when using the intermittent (as-needed) schedule most appropriate to the surge context, can continue indefinitely, whereas IPA assignments would require renewal of agreements at the two-year mark (and could run into statutorily and regulatorily mandated time limits at the four-year mark).

[92] OPM, "Fact Sheet: Expert and Consultant Pay." For 2025, the base salary for GS-15, step 10, is approximately $163,000 per year. See OPM, "Salary Table 2025-GS."
[93] 5 C.F.R. 304.103(b)(3). The Code of Federal Regulations also specifies that experts/consultants cannot make final decisions on substantive policies or otherwise function in the agency chain of command.
[94] 5 C.F.R. 304.103(b)(4)-(5).
[95] See 5 C.F.R. 304.102(g).
[96] Specifically, experts/consultants who work "on a strictly intermittent basis" (i.e., no regular work schedule) may be appointed without any time limit, or for any period determined by the agency. 5 C.F.R. 304.103(a)(2). All others must be given an initial appointment not to exceed one year, see 5 C.F.R. 304.103(a)(2); 5 C.F.R. 304.102(g), and those who work no more than 1040 hours per year (i.e., half-time) can continue to receive successive one-year appointments, see 5 C.F.R. 304.103(c)(2)(i). If working more than 1040 hours per year, however, experts/consultants can receive only a one-year extension of their initial one-year appointment, for a total service period of two years. See 5 C.F.R. 304.103(c)(1). OPM may grant exceptions to these time limits in "unforeseen and unusual circumstances." 5 C.F.R. 304.103(c)(3).

Subsection 304.103(c)(2)(ii) outlines a separate reappointment option, setting a lifetime limit for expert/consultant pay and assessing an individual's cumulative earnings against that standard. These two paths to reappointment are mutually exclusive. See 5 C.F.R. 304.103(c)(2).
[97] 5 C.F.R. 304.103(a)(2).
[98] OPM, "Fact Sheet: Expert and Consultant Pay."

A necessary caveat, however, is that experts/consultants are barred by regulation from doing work performed by the agency's regular employees. That means that tiers with talents pools that encompass teams at federal agencies (e.g., Engineering & Hardware Implementation & Monitoring) may have a harder time utilizing the expert/consultant authority to bring in outside AI security talent, though that will of course depend on the exact details of the work envisioned for the outside experts, in comparison to the work performed by regular agency employees. Agencies' internal teams conducting regulatorily required reviews of potential expert/consultant appointments are attuned to this limit, not least because agencies' OIGs are known to investigate whether they are using special appointment authorities appropriately.[99] As such, the scope of expert/consultant positions could get narrowed for this reason during an agency's required internal review process, or—if narrowing isn't feasible—reviewers might even deny use of the expert/consultant authority.

While expert/consultant appointments lack the cost-sharing flexibility of IPAs, **the government does have authority to accept expert/consultant services on an unpaid basis**, if individuals are willing to waive their right to compensation. Though that might seem unrealistic, it is worth noting that unpaid experts/consultants are exempted from the general prohibition against accepting compensation for services to the government from anyone outside of the government.[100] Thus, unpaid experts and consultants could continue to collect their regular salaries from their outside employers—which, as with IPAs, those outside employers might be willing to continue paying given the intermittent, short-term appointments contemplated in this project.

That said, expert/consultant appointments share some of the same time-to-hire challenges as IPA assignments, given regulatory requirements for internal agency review processes, which could include entering into the aforementioned OIG-recommended agreements for each expert/consultant appointment. While streamlining those processes to a degree is probably possible in a crisis, some baseline level of review (particularly of the duties and responsibilities of positions) will still be necessary for agencies to meet their regulatory obligation for a system of controls and oversight. Even with a streamlined process, bottlenecks could emerge if there's demand for many new positions at once. Therefore, as with IPAs, **advance planning—including steps to hire, clear, and onboard individuals on as-needed status so that they are on standby and ready to serve immediately in the event of a crisis—remains the wisest course of action.**

---

[99] See, e.g., Office of Inspector General, Federal Trade Commission, "Audit of the Federal Trade Commission's Unpaid Consultant and Expert Program," 17 (in response to OIG findings, FTC management committing to reviewing, updating, and finalizing internal policy manual to "provide guidance on allowable vs prohibited activities" for experts/consultants).
[100] 18 U.S.C. 209(c).

Relatedly, agencies have tended to use the expert/consultant authority somewhat sparingly,[101] which could exacerbate the strain on agencies' internal review processes during a rapid and extensive scale-up for crisis response. Though there is generally no limit on the number of expert/consultant appointments an agency can make, utilizing this authority in service of surge hiring could necessitate not only a rethink of internal processes but also a shift in mindset about the scope of this authority.

Schedule A Appointments

Background: When specifically authorized by OPM, **agencies may make appointments under a regulatory framework called Schedule A to positions that are not "confidential or policy-determining" in character, and for which it is not practicable to examine**, as would ordinarily happen with competitive-service hiring.[102] Schedule A's existing government-wide authorizations cover a broad array of uses, ranging from appointments of chaplains, attorneys, and individuals with disabilities to positions filled by non-U.S. citizens or nationals in the absence of qualified candidates.[103]

Among the uses of Schedule A already authorized by OPM as available to all federal agencies,[104] the following are most relevant for surge hiring AI security talent:

- Subsection 213.3102(i)(2) authorizes hiring for temporary and less-than-full-time positions "for which a **critical hiring need** exists." Such appointments are limited to 30 days, with the possibility of a one-time extension for an additional 30 days "if continued employment is essential to the agency's operations."[105]

---

[101] See, e.g., Office of Inspector General, Federal Trade Commission, "Audit of the Federal Trade Commission's Unpaid Consultant and Expert Program," 3 (noting that the FTC had nine unpaid expert/consultant positions in Fiscal Year (FY) 2021, versus two such positions in FYs 2019-2020).
[102] 5 C.F.R. 213.3101; see also 5 U.S.C. 3302. "Confidential or policy-determining character" previously was defined as meaning "of a character exclusively associated with a noncareer political appointment that is identified by its close working relationship with the President, head of an agency, or other key appointed officials . . . ," though regulations containing such definitions have been stayed indefinitely. See Executive Order 14171 (signed January 20, 2025).
[103] See generally 5 C.F.R. 213.3102. One of these aforementioned authorities, 5 C.F.R. 213.3102(bb), allows for hiring of non-U.S. citizens under certain circumstances, and so might be a fruitful line of inquiry for future work on surging AI security talent, given that a significant portion of AI talent already studying or working in the United States does not have U.S. citizenship. See, e.g., The White House, "AI Talent Report" (noting that 59% of AI-related PhDs awarded in 2022 by U.S. institutions were to non-U.S. citizens). Nonetheless, given the challenges in hiring non-U.S. citizens for federal government work, particularly with obtaining clearances, this category of talent is beyond the scope of this report.
[104] OPM also makes agency-specific Schedule A authorizations, though given their more limited applicability, this report does not address them. See, e.g., Federal Register, "Excepted Service; Consolidated Listing of Schedules A, B, and C Exceptions," 89 Fed. Reg. 61177 (providing current list of agency-specific authorities).
[105] No individual can be employed under this authority for more than 60 days in any 12-month period. 5 C.F.R. 213.3102(i)(2).

- Subsection 213.3102(*l*) authorizes hiring for "[p]ositions requiring the temporary and intermittent employment of **professional, scientific, or technical experts** for consultation purposes."[106]
- Subsection 213.3102(o) authorizes hiring for "[p]ositions of a **scientific, professional, or analytical nature when filled by bona fide members of the faculty of an accredited college or university** who have special qualifications for the positions to which appointed." Such appointments cannot exceed 130 working days per year.
- Subsection 213.3102(r) authorizes hiring for "positions established in support of fellowship and similar programs that are filled from limited applicant pools and operate under specific criteria developed by the employing agency and/or a non-Federal organization." Such programs include "**professional/industry exchange programs** that provide a cross-fertilization between the agency and the private sector to foster mutual understanding, an exchange of ideas, or to bring experienced practitioners to the agency."[107]

When it comes to recruitment and time-to-hire, Schedule A's wide range makes generalization and estimation challenging. **While Schedule A hiring is expedited as compared to normal competitive-service hiring because procedures like examination are not required, average timelines still generally run closer to weeks and months**, rather than days.[108] Some situations might require more lead time—for example, if an agency is setting up a new exchange program pursuant to subsection 213.3102(r).

In contrast to IPA assignments and expert/consultant appointments, **Schedule A appointments do not come with a cost-sharing mechanism or unpaid option** that would allow the government to defray the cost for salaries and other expenses.[109] Individuals hired

---

[106] Consistent with time limits outlined elsewhere in the Schedule A regulations, a position under the authority in subsection 213.3102(*l*) could be extended indefinitely if the individual works intermittently for a total of less than six months (or 1040 hours) per year. 5 C.F.R. 213.104(b)(3)(i). Otherwise, an appointment under this position is capped at two years, see 5 C.F.R. 213.104(b)(1), with a bar on refilling a position when it has been occupied for an aggregate of two years in the preceding three-year period, see 5 C.F.R. 213.104(b)(2).

[107] Appointments under this authority are limited to four years. 5 C.F.R. 213.3102(r).

[108] See Partnership for Public Service, "IPA Agency and Candidate Guidebook," 23 (commenting that the timelines for Schedule A hiring likely correspond to those for IPA hiring, i.e., approximately three to six months).

[109] In general, the federal government is prohibited by the Antideficiency Act from accepting unpaid volunteer services unless a particular government-wide or agency-specific authorization exists. See OPM, USAJobs, "Volunteers"; 31 U.S.C. 1342. Per OPM, these few generally applicable exceptions are student volunteers, emergencies involving the protection of life or property, assistants to handicapped employees, and–as previously discussed–the expert/consultant authority. See OPM, USAJobs, "Volunteers."

Schedule A does include a subsection 213.3102(k) regarding "[p]ositions without compensation provided appointments thereto meet the requirements of applicable laws relating to compensation." But given the reference to otherwise-applicable compensation laws, it seems doubtful that subsection 213.3102(k) could be used on its own as a way to hire AI security talent on an unpaid basis, without some other specific authorization waiving the Antideficiency Act's prohibition and/or other laws.

under Schedule A are generally paid according to the General Schedule (GS) pay system that applies to competitive-service employees.[110]

Limits on the type of work to be performed and the total allowable time in service vary depending on the specific Schedule A authority at issue. Regarding the type of work, some specify that it be scientific, technical, or analytical in its nature or focus (i.e., subsections 213.3102(*l*) and (o)), whereas another requires only any kind of "critical hiring need" (i.e., subsection 213.3102(i)(2)). Regarding time limits, some are quite short (e.g., subsection 213.3102(i)(2)'s cap of 60 days per year), while others could more easily be extended (e.g., subsection 213.3102(r)'s upper limit of four years). Unless otherwise specified, Schedule A appointments may be full-time, part-time, or intermittent,[111] though—as with the expert/consultant authority—the work schedule chosen can have knock-on effects for an individual's total allowable time in their federal appointment. And aside from subsection 213.3102(o), which requires that appointees be "bona fide members of the faculty of an accredited college or university," the Schedule A authorities surveyed here provide flexibility regarding an individual's current employer or affiliation, meaning it can generally be used to hire individuals from the private sector.[112]

Individuals appointed under Schedule A are federal employees, and so are subject to the generally applicable ethics and conflicts of interest rules, as well as the standard post-government restrictions.

Application: Schedule A is perhaps best thought of as a collection of multiple authorities, rather than a single authority, and that variety is both an advantage and disadvantage relative to the IPA and expert/consultant authorities. Notably, **OPM has the power to specifically authorize *additional* Schedule A authorities, either on a government-wide or agency-specific basis, without the need for Congressional sign-off.** Thus, to the extent that the Schedule A authorities described in this report might fail to provide comprehensive coverage in the event of an AI-related crisis, OPM can authorize more AI-crisis-specific uses of Schedule A. That regulatory flexibility amounts to a significant edge over the IPA and expert/consultant authorities, which are largely creatures of statute, meaning that action by Congress would generally be required to change their terms.

---

[110] Shimabukuro and Staman, "Categories of Federal Civil Service Employment: A Snapshot."
[111] 5 C.F.R. 213.102(b)(3)(i).
[112] Subsection 213.3102(r) could involve certain limitations relating to the current employer or affiliation of prospective hires–e.g., individuals participating in "professional/industry exchange programs" would likely have to come from the private sector–though that authority gives agencies considerable latitude in establishing criteria for such programs and hiring.

"Schedule A authorities in general have relatively few required steps and processes."

Another advantage is that most—though not all—of the Schedule A authorities detailed above can be used to hire individuals from anywhere, i.e., with no limits stemming from the type of organization with which they are currently employed.[113] That flexibility positions Schedule A favorably in relation to the IPA authority, which cannot be used for individuals coming from the private sector, though somewhat less favorably in comparison to the expert/consultant authority, which contains no restrictions whatsoever regarding individuals' current employers or affiliations.

As with the IPA and expert/consultant authorities, there is no blanket limit on the number of appointments that may be made under the Schedule A authorities discussed here.[114] However, as compared to the agreements and eligibility certifications necessary for IPAs, as well as the regulatorily mandated internal review protocols for expert/consultant appointments, **Schedule A authorities in general have relatively few required steps and processes.** Thus, particularly with respect to an authority like subsection 213.3102(i)(2), which is keyed to extremely short (i.e., 30-60 day) appointments to address critical hiring needs, it is conceivable that Schedule A could be deployed during a crisis in just days or a couple of weeks—though, for reasons like recruitment and obtaining security clearances, advance planning and action are still prudent.

On the other hand, Schedule A lacks the cost-sharing and unpaid options that the IPA and expert/consultant authorities offer, respectively, which may make Schedule A comparatively costly for the government. That said, individuals appointed under Schedule A would receive government pay while potentially also continuing to receive their normal outside salaries (which is possible if they receive SGE designation, discussed later in this section, and if outside companies agree to keep paying their employees during periods of government service). While this may not matter much to the most highly compensated AI security talent working at frontier companies and elsewhere, it could have the upside of attracting more AI security talent from academia, nonprofits, and smaller companies, as well as independent researchers, given the prospect of extra earnings during individuals' time spent in government service.

[113] The exception being the Schedule A authority contained in 5 C.F.R. 213.3102(o), which requires that appointees be on the faculty of an accredited college or university.
[114] It is worth noting, however, that many of the agency-specific Schedule A authorities do have such caps. See Federal Register, "Excepted Service; Consolidated Listing of Schedules A, B, and C Exceptions," 89 Fed. Reg. 61177.

Turning now to the specific Schedule A authorities outlined above, subsection 213.3102(i)(2)'s **"critical hiring need" authority is both broad in scope and particularly suited to the crisis scenarios envisioned here**, in which there is acute need for AI security talent. Additionally, this authority's focus on critical hiring need is less restrictive than the expert/consultant authority's prohibition on being used to do work performed by the agency's regular employees, since "critical hiring need" also encompasses work that an agency might already be doing, just not in the numbers needed. Thus, this authority might have fewer issues reaching those areas of the Typology of Talent (e.g., Engineering & Hardware Implementation & Monitoring) for which the talent pools are already comprised of various federal agencies and subcomponents. That said, this authority is sharply time-limited (i.e., no more than 60 days in any 12-month period), leaving little flexibility if more time ends up being needed for training and/or crisis response.

Subsection 213.3102(o), which **authorizes hiring for scientific positions when filled by members of the faculty of an accredited college or universit***y*, has much in common with the IPA authority, which also draws from academia. While subsection 213.3102(o)'s scope is somewhat narrower (notably, it does not encompass nonprofits, as the IPA authority does), it does not require agreements with lending institutions, which could well expedite the time-to-hire, particularly if an agency begins hiring only after a crisis emerges, without having done the advance planning advocated for here. Subsection 213.3102(o)'s time limit of 130 working days per year also accords with the short-term surge hiring needed for the specific project of crisis response, and aligns well with the time limits for SGE designation, which mitigates somewhat the barrier to hiring posed by conflicts of interest restrictions.

Finally, while subsection 213.3102(*l*)'s **authority to hire scientific and technical experts** is quite similar to the expert/consultant authority contained in 5 U.S.C. 3109, its lack of any restriction on such experts doing work that is also performed by the agency's regular employees means that it could potentially be deployed to cover a broader array of the Typology of Talent than the section 3109 expert/consultant authority alone. As for subsection 213.3102(r)'s **authority to hire positions related to fellowships and similar programs**, including "professional/industry exchange programs," it could potentially be used to develop crisis-response teams drawing from particular companies or other parts of the private sector to which the government has few pipelines or existing connections, though considerable coordination ahead of any crisis would be needed to stand up such programs.

Non-Hiring Mechanisms

Special Government Employees

As one possible way to boost surge-hiring efforts, designating AI security talent as Special Government Employees (SGEs) could potentially reduce the time-to-hire and related barriers to government service that stem from conflicts of interest rules. In contrast to IPA assignments and appointments under the expert/consultant and Schedule A authorities, "SGE" is not itself a hiring authority or mechanism, but rather a status or designation that can be conferred on federal employees once hired. For example, an expert or consultant hired under 5 U.S.C. 3109 might be designated as an SGE, taking them out of the default category of "regular" government employees—the upshot being that SGEs are subject to "somewhat less restrictive" conflicts of interest requirements relative to other federal employees.[115]

Federal conflicts of interest rules are considerable, with potential conflicts arising from, among other things, equity-related interests that a federal employee continues to hold in a company for which the employee works or worked.[116] Common equity-related interests include stock, incentive stock options, restricted stock, restricted stock units (RSUs), stock appreciation rights, phantom stock, warrants, and employee stock option plans (ESOPs) and employee stock purchase plans (ESPPs).[117] Equity-related interests like RSUs often comprise a portion of employee compensation packages at tech companies.[118] As the Office of Government Ethics (OGE) advises, the potential for conflicts of interest arises as soon as an equity-related interest is received, even before any vesting has occurred or any option has been exercised.[119] Thus, full application of federal conflicts of interest restrictions to AI security talent could frustrate surge hiring during a fast-unfolding AI-related crisis.

"Short-term crisis-response appointments are a uniquely good match for SGE designation."

As a first step, the agency employing an individual (and not OGE or OPM) bears responsibility for determining whether they qualify for SGE status.[120] By law, **SGEs are limited to performing temporary duties either on a full-time or intermittent basis for not more than 130 days during any 365-day period.**[121] To designate a prospective employee as an SGE, the agency must make a "good faith estimate," prior to that employee beginning work, that they are expected to serve no more than 130 days per year, so that both the employee and agency are aware of the applicable restrictions from the outset.[122] While this 130-day limit

[115] Office of Government Ethics, "Summary of the ethics provisions that apply to Special Government Employees (SGEs)," 1; see also 18 U.S.C. 202(a).
[116] Office of Government Ethics, "Conflicts of Interest Considerations: Corporate Employment," 3.
[117] Office of Government Ethics, "Conflicts of Interest Considerations: Corporate Employment," 3-4.
[118] See, e.g., Founders Network, "A Founder's Ultimate Guide to Startup Equity."
[119] Office of Government Ethics, "Conflicts of Interest Considerations: Corporate Employment," 4.
[120] Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees," 1.
[121] 18 U.S.C. 202(a).
[122] Office of Government Ethics, "Memorandum dated February 15, 2000, from Stephen D. Potts, Director, to Designated Agency Ethics Officials, General Counsels and Inspectors General Regarding Summary of Ethical Requirements Applicable to Special Government Employees," 5. Per that OGE memorandum, if an employee "unexpectedly serves more than 130 days during the ensuing 365-day

is a nonstarter when individuals are hired for longer-term appointments to address longer-term projects, **short-term crisis-response appointments like those envisioned here are a uniquely good match for SGE designation.** Furthermore, there is no limit on the number of employees an agency may designate as SGEs.[123]

While SGEs have less restrictive conflicts requirements relative to regular government employees, **SGEs are nonetheless government employees and so "subject to many more requirements than non-employees," such as independent contractors to whom such rules generally do not apply.[124]** Furthermore, SGEs often have "substantial outside activities and financial interests that may raise difficult ethics questions."[125] Thus, while SGE designation may make government service more attractive to AI security talent, it may not always reduce time-to-hire or the burden on agency personnel involved in hiring, including agencies' ethics offices.

In contrast to nearly all regular government employees, SGEs are exempted from the prohibition against accepting compensation for services to the government from anyone outside of the government.[126] In practice, that exemption means that SGEs "may continue to collect their regular salary from an outside employer for days on which they are providing services to the Government," even when their government service is paid.[127] That would undoubtedly increase the appeal of government service to AI security talent, in that they need not necessarily forgo any part of their normal salary while serving. Similarly, accelerated vesting of RSUs and other equity-related interests, which might be part of divesting to remedy a conflict of interest that arises, could run afoul of the prohibition against outside compensation for government services if not for SGE status.[128] Additionally, of particular interest to IPA assignees and others hailing from nonprofits or academia, SGEs are subject to a narrower set of restrictions on when they can receive outside compensation for teaching, speaking, or

---

period, the individual still will be deemed an SGE for the remainder of that period," although the agency must reevaluate whether SGE designation continues to be appropriate for the next 365-day period.

[123] See U.S. Government Accountability Office, "Federal Workforce: Opportunities Exist to Improve Data on Selected Groups of Special Government Employees," 10. In its report, GAO found that agencies use SGE designation most frequently for individuals serving on federal boards, though about three percent of SGEs were not serving on federal boards and were employed as experts/consultants. Between 2005 and 2014, the annual average number of SGEs across federal agencies was approximately 2,000, not counting those SGEs serving on federal boards.

[124] Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees," 1.

[125] Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees," 1.

[126] See 18 U.S.C. 209(c). For current purposes, the other notable exemption from this rule is *unpaid* experts/consultants hired pursuant to 5 U.S.C. 3109, by virtue of their uncompensated status.

[127] Office of Government Ethics, "Memorandum dated February 15, 2000, from Stephen D. Potts, Director, to Designated Agency Ethics Officials, General Counsels and Inspectors General Regarding Summary of Ethical Requirements Applicable to Special Government Employees," 17.

[128] Office of Government Ethics, "Conflicts of Interest Considerations: Corporate Employment," 4-5.

writing.[129] Such latitude could make government service more attractive to AI security talent coming from research institutions.

That said, there are some core conflicts of interest rules from which SGEs receive no exemption or variance. Namely, **all employees, including SGEs, are prohibited by federal criminal law from personally and substantially participating in any particular matter that would have a direct and predictable effect on the employee's financial interests or those imputed to the employee** (e.g., those of a spouse or an outside entity with whom the individual is employed or negotiating for employment).[130] Likewise, SGEs and all other employees are permanently restricted under criminal law from representing anyone before the government on a particular matter involving specific parties (e.g., a research program) in which the employee participated personally and substantially while working for the government, and in which the government is a party or has a substantial interest.[131]

While waivers from these generally applicable restrictions might be possible in certain circumstances,[132] the minimum set of conflicts rules to which SGEs are subject remains significant. Thus, future work on building government surge capacity for AI-related crises should examine in more depth tools beyond federal hiring for obtaining AI security talent and expertise, such as public-private partnerships and—as addressed compactly below—use of contractors.

Contractors

Contracting is likely to be another useful tool for surging AI security talent. **Contracting has several advantages—most notably, the flexibility it provides the government in obtaining services for a specific purpose while reducing the overhead and long-term commitment that comes with direct hiring.** This is all the more advantageous when much of the talent the government needs is not interested in leaving the private sector for long-term government service. **Contracting could also prove a faster means of surging AI security talent in some instances, particularly if agencies work ahead of time to establish agreements before a crisis occurs.**

---

[129] Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees," 8; see also Office of Government Ethics, "Memorandum dated February 15, 2000, from Stephen D. Potts, Director, to Designated Agency Ethics Officials, General Counsels and Inspectors General Regarding Summary of Ethical Requirements Applicable to Special Government Employees," 20-22.
[130] 18 U.S.C. 208(a); see also Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees," 3.
[131] 18 U.S.C. 207(a)(1); see also Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees," 10; MIT Office of the General Counsel, "MIT Employees Considering Service as a Special Government Employee." MIT's General Counsel's office has—perhaps in an abundance of caution—intepreted this to mean that "personal and substantial involvement in a matter as an SGE will preclude involvement in the same matter upon return to MIT," demonstrating the potential chilling effect of such restrictions.
[132] See generally Office of Government Ethics, "Ethics Laws Applicable to Special Government Employees."

It is no surprise, then, that contracting has previously been used to provide "surge" capacity, including within the Intelligence Community,[133] and is more regularly used to provide "supplemental expertise in specialized fields"—from linguistics to weapon systems maintenance—on a more continuing basis.[134]

"Contracting could prove a faster means for surging AI security talent if agencies establish agreements before a crisis occurs."

Contracting could be used to facilitate a surge in several ways:

1. Utilizing **acquisition flexibilities in the Federal Acquisition Regulation (FAR)** when (and potentially before) a crisis emerges;
2. Redirecting **existing contracts that touch on AI security talent and expertise** to response efforts when a crisis emerges;
3. Establishing **new indefinite delivery, indefinite quantity (IDIQ) contracts** for AI security talent ahead of a crisis and then placing task orders for services pursuant to those contracts as needed, once a crisis occurs; and
4. Establishing **Other Transaction Agreements (OTAs)** aimed at crisis response (ideally in place before a crisis occurs).

First, the FAR—an extensive body of rules that governs how the federal government procures goods and services—contains some "acquisition flexibilities" that can speed up the standard and often time-consuming contracting process.[135] Most notably, **an agency can depart from the default "full and open" competition procedures when a contracting officer determines that there is "unusual and compelling urgency."[136]** In that situation, the government is exempted from public posting of a synopsis of the solicitation, and normal contractor registration requirements are waived.[137] Even more significantly, contracts based on these circumstances are not held, suspended, or terminated when a bid protest is filed within a certain time period, as otherwise happens pending GAO's decision on a timely protest.[138] The

---

[133] Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 8; see also Neenan, "Defense Primer: Department of Defense Contractors." As one data point, a government-run inventory of the IC's contractor workforce for Fiscal Year 2007 showed that "5% supported surge requirements." Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 9.
[134] Neenan, "Defense Primer: Department of Defense Contractors."
[135] See Fiorentino, "Federal Procurement: Emergency-Related Acquisition Flexibilities," 4. Notably, some of the FAR's "acquisition flexibilities" (including the ones detailed here) can be used even when there is not a formal emergency declaration in place. See Fiorentino, "Federal Procurement: Emergency-Related Acquisition Flexibilities," 4.
[136] Fiorentino, "Federal Procurement: Emergency-Related Acquisition Flexibilities," 4-5.
[137] See Fiorentino, "Federal Procurement: Emergency-Related Acquisition Flexibilities," 5.
[138] See Fiorentino, "Federal Procurement: Emergency-Related Acquisition Flexibilities," 5.

"unusual and compelling urgency" flexibility thus could be quite valuable during an AI-related crisis, though more work is needed to assess whether and to what extent other non-waived requirements of the FAR might still impede its usage in the context of surging AI security talent.

Second, **agencies might be able to simply redirect existing contracts for AI-related services toward crisis response, saving the time and expense associated with getting new contracts into place.** Specifically, one expert consulted for this project noted that many government contracts, especially those with the IC, include provisions that allow an agency to change the scope or focus of work performed under a contract if the mission demands it. More broadly, the General Services Administration (GSA), which provides centralized procurement options for the federal government, has in place a Governmentwide Acquisition Contract (GWAC) called Alliant 2 that covers AI and other emerging technologies, and includes vendors such as IBM.[139] GWACs like Alliant 2 are multiple-award contracts, meaning that they are awarded to multiple vendors who then compete among themselves for task orders that the government issues pursuant to the contract.[140] GWACs are also indefinite delivery, indefinite quantity (IDIQ) contracts, meaning that—unlike conventional contracts—they do not specify the precise quantities and delivery timelines for goods or services (beyond a preset minimum quantity), with the government instead placing orders only as needed.[141] While Alliant 2 and other existing contracts might be of some use to agencies in responding to an AI-related crisis, their apparent focus on application-level services (supporting agencies' efforts "to modernize their IT infrastructure, strengthen cybersecurity capabilities, and drive innovation across their operations") seems aimed more at improving day-to-day performance, and potentially less relevant to the core competencies that the Typology of Talent identified as critical to addressing an AI-related crisis.[142] Thus, simply redirecting existing contracts is unlikely to be sufficient, meaning that agencies will need to establish new contracts to address the full range of crises presented in Section 1.

The likely inadequacy of these existing contracts points to a third option: establishing new IDIQ contracts for AI security talent. Ideally, agencies would establish these IDIQ contracts in advance—before a crisis occurs—and then be able to quickly place task orders for services when a crisis begins. This approach has several benefits. For starters, it is a well-precedented strategy. The FAR specifically notes that multiple-award IDIQs are the government's preferred contracting method,[143] and FEMA has long used indefinite delivery contracts as part of its "advance contract strategy" for emergency response, "preclud[ing] the need to procure goods and services under [the] unusual and compelling urgency" mechanism

[139] U.S. General Services Administration, "Alliant 2 Governmentwide Acquisition Contract (GWAC)"; U.S. General Services Administration, "Alliant 2 industry partners."
[140] U.S. General Services Administration, "Alliant 2 Governmentwide Acquisition Contract (GWAC)"; see also U.S. General Services Administration, "Ways you can sell to government."
[141] U.S. General Services Administration, "Alliant 2 Governmentwide Acquisition Contract (GWAC)"; see also Fiorentino and Neenan, "Indefinite Delivery, Indefinite Quantity Contracts."
[142] U.S. General Services Administration, "Alliant 2 Governmentwide Acquisition Contract (GWAC)."
[143] See Fiorentino and Neenan, "Indefinite Delivery, Indefinite Quantity Contracts."

in the FAR.[144] **IDIQs are also fast where it matters.** While establishing a new IDIQ contract generally takes several months or more, once in place, IDIQs can "speed service delivery"[145] because—according to the estimates of experts interviewed for this report—task orders pursuant to an existing IDIQ contract can be issued and finalized within a matter of weeks (or at the upper end, one to two months). Multiple-award IDIQs can also lessen the risk of "vendor lock" and "keep[] pressure on pricing" because multiple vendors continue to compete to win the task orders issued under the contract,[146] though that level of competition is not guaranteed where several of the relevant talent pools are dominated by only a handful of companies, as is the case here.[147] All things considered, IDIQs are likely to be a valuable tool for surging AI talent if established before a crisis and set up with adequate competition.

Fourth, **authorized agencies could use OTAs, which offer some benefits over traditional government procurement contracts.[148]** Notably, the government made significant use of OTAs for COVID-19 response efforts, providing a precedent for using this instrument in the context of crisis response.[149] **Most importantly, OTAs are not subject to the FAR and certain procurement statutes**, giving agencies considerable latitude in deciding how to structure agreements and choose participants,[150] though statutes and agency guidelines often still emphasize that competitive procedures be used to the maximum extent practicable.[151] OTAs can also involve cost-sharing between the government and participants, providing agencies with yet more flexibility.[152] Furthermore, agencies can award OTAs to individual organizations or to "consortia," which can be composed of traditional defense or other contractors, "nontraditional" companies (i.e., those that generally do not contract with the government), nonprofits, and academic institutions.[153]

---

144 U.S. Government Accountability Office, "Covid-19 Contracting: Observations on Federal Contracting in Response to the Pandemic," 8-9.
145 Fiorentino and Neenan, "Indefinite Delivery, Indefinite Quantity Contracts" (contrasting the speed of service delivery on IDIQ task orders with the "traditionally [] long procurement lead time to award" when initially establishing IDIQ contracts).
146 Fiorentino and Neenan, "Indefinite Delivery, Indefinite Quantity Contracts."
147 Fiorentino and Neenan, "Indefinite Delivery, Indefinite Quantity Contracts."
148 See Halchin, "Other Transaction (OT) Authority," 1 (observing that OTAs are "defined in the negative," that is, they are "not a contract, grant, or cooperative agreement," but rather a "special type of vehicle or instrument" for entering into an agreement with the government).
149 See U.S. Government Accountability Office, "Innovation Lab"; see also Komitsky, "OTA Offers Big Benefits . . . If You Know How to Use It."
150 Halchin, "Other Transaction (OT) Authority," 1-2.
151 See Vadiee and Garland, "The Federal Government's 'Other Transaction' Authority," 12 (quoting DOD, Other Transactions Guide for Prototype Projects §§ C1.2.3, C2.1.3.1.6, January 2017).
152 See Advanced Technology International (ATI), "Other Transaction Agreements (OTA)."
153 See U.S. Government Accountability Office, "Other Transaction Agreements: DOD Can Improve Planning for Consortia Awards," 1. Awarding OTAs to consortia can provide similar benefits as multiple-award IDIQs, insofar as members of the consortia compete for project sub-awards and so may keep prices down. See U.S. Government Accountability Office, "Other Transaction Agreements: DOD Can Improve Planning for Consortia Awards," 27.

However, OTAs come with some limitations. For starters, Congress has thus far granted OTA authority to only a handful of agencies, including DOD, the Department of Energy, and DHS.[154] Most others, including the Department of Commerce and the Office of Science and Technology Policy, lack such authority. Furthermore, in authorizing provisions, OTAs are described as agreements "other than contracts, cooperative agreements, and grants,"[155] which indicates that they should not be used when, for example, the "principal purpose" of the agreement is "to acquire (by purchase, lease, or barter) property or services for the direct benefit or use of the United States Government."[156] Rather, as outlined in various statutory provisions as well as regulations and agency guidance, OTAs are generally aimed at research projects and developing prototypes, with an emphasis on the novelty of the research or application.[157] Thus, to the extent that an OTA for AI-crisis response might principally serve to acquire goods or services, it may not be the most appropriate vehicle to use.[158]

By contrast, OTAs in the context of the COVID-19 response centered on efforts like vaccine development, as well as developing new therapeutics and diagnostic tools, involving the sort of innovations for which OTAs are intended.[159] While it is conceivable that plans for AI crisis response could involve similarly innovative efforts (e.g., private companies partner with government in building a hardened system or tech stack for crisis response and/or subsequent usage, relying on newly developed and/or novelly applied technology), **it is unlikely OTAs could be used simply to surge AI security talent into government without some sort of new research or application hook.**

Additionally, while "[i]t appears that a favorable consensus exists regarding the use of OT authority . . . based largely on the experiences and observations of individuals who participate, or have participated, in OT[A]s," it is uniquely difficult to quantitatively evaluate the benefits of OTAs.[160] That is at least in part because—unlike with other government contracts—the ordinary audit services and quality assurance measures from the FAR do not apply.[161] Finally, as GAO has noted, OTAs actually "may take longer to award than FAR-based contracts due to factors

---

[154] See Vadiee and Garland, "The Federal Government's 'Other Transaction' Authority," 1; see also Komitsky, "OTA Offers Big Benefits . . . If You Know How to Use It."
[155] E.g., 10 U.S.C. 4021(a).
[156] 31 U.S.C. 6303(1).
[157] See Halchin, "Other Transaction (OT) Authority," 1; see, e.g., 10 U.S.C. 4021(a) (granting DOD OTA authority "in carrying out basic, applied, and advanced research projects") & 4022(e) (in granting DOD OTA authority for prototypes, defining a "prototype project" as, among other things, "a pilot or novel application of commercial technologies for defense purposes").
[158] See, e.g., Vadiee and Garland, "The Federal Government's 'Other Transaction' Authority," 6 (explaining that DHS uses research OTAs for "transactions for which the government's acquisition of goods and services is not the principal purpose" (quoting DHS, Management Directive No. 0771.1, Other Transaction Authority 4 (July 8, 2005))).
[159] See Komitsky, "OTA Offers Big Benefits . . . If You Know How to Use It."
[160] Halchin, "Other Transaction (OT) Authority," 28.
[161] See Halchin, "Other Transaction (OT) Authority," 23.

such as drafting and negotiating all the terms and conditions," given how much flexibility is possible, placing a premium on having any AI-specific OTAs in place pre-crisis to be effective.[162]

Barriers to contracting: As with direct federal hiring authorities, **speed is the biggest barrier to surging AI security talent during a crisis via contracting.** Luckily, agencies can accelerate how quickly services and other solutions are provided via contracting through advance planning—namely, by establishing relevant IDIQs and OTAs ahead of a crisis.[163] Such advance planning is critical given that, under the FAR, selecting a vendor can sometimes take up to two years, meaning that "technologies that are considered state-of-the-art when a new procurement is envisioned are often outdated by the time a contract is awarded."[164]

"Agencies can accelerate how quickly services and other solutions are provided via contracting through advance planning—namely, by establishing relevant IDIQs and OTAs ahead of a crisis."

Another potential barrier, also stemming from the FAR, has to do with what work contractors can perform. Specifically, agencies must ensure (via a written determination accompanying the statement of work)[165] that contractors do not perform "inherently governmental activities," which are reserved to direct federal employees and officials.[166] Inherently governmental activities include functions that involve exercising discretion or making value judgments on behalf of the government, though the term generally does not extend to merely gathering information or providing advice, opinions, or ideas to government officials to inform their own decisionmaking processes.[167] As described in Section 1, the work that AI security talent would most likely perform in response to an AI-related crisis includes diagnosing the root causes of AI system dysfunction, rapid assessment of AI capabilities to predict any further attacks, implementing rapid containment measures, and deploying remedial fixes to halt further harm. **That advisory and implementation-focused work likely does not entail the sort of discretion that**

[162] U.S. Government Accountability Office, "Defense Acquisitions: DOD's Use of Other Transactions for Prototype Projects Has Increased," 8.
[163] U.S. Government Accountability Office, "Covid-19 Contracting: Observations on Federal Contracting in Response to the Pandemic," 16-17 (noting that "agencies can leverage contracts awarded in advance of a disaster to rapidly and cost-effectively mobilize resources," which "can help preclude the need to procure critical goods and services noncompetitively").
[164] Vadiee and Garland, "The Federal Government's 'Other Transaction' Authority," 2 (quoting CIO Council, IT Acquisition and Contracts Management, 2016 State of Federal IT Report, at F-2).
[165] See Federal Acquisition Regulation (FAR), "Subpart 7.5 - Inherently Governmental Functions," 7.503(e).
[166] Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," ii (Summary).
[167] See Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 3-4 & n.9, 22 (Appendix A).

necessarily makes it an inherently governmental activity, though it could qualify as a "critical" function, or as "closely associated" with inherently governmental functions, particularly given its potential effects on high-stakes national-security outcomes.[168]

While critical and closely associated functions can be performed by contractors,[169] the Office of Management and Budget (OMB) has advised that agencies must put in place certain safeguards so that contractors do not become the de facto decisionmakers, their recommendations accepted and implemented unquestioningly.[170] For example, agencies must have "sufficient internal capability"—i.e., federal employees with the necessary experience and expertise—to evaluate contractors' work product and "formulate alternatives,"[171] and "establish[] in advance a process for subjecting the contractor's discretionary decisions and conduct to meaningful oversight and, whenever necessary, final approval by an agency official."[172] To this end, contractors are barred from directing or otherwise controlling federal employees,[173] and one contractor's evaluation of a separate contractor's performance—while not inherently governmental—is a closely associated function that must be appropriately monitored.[174] For these reasons as well, **advance planning before a crisis is vital to ensure agencies have the capacity needed, both in terms of oversight processes and their own in-house expertise, to appropriately manage contractors' activities.** Among other things, agencies must determine whether (and if so, how) they can place any surged contractor talent in relation

---

[168] See Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 3-4. A "critical" function is "necessary to the agency being able to effectively perform and maintain control of its mission and operations," and is often (though not necessarily) "recurring and long-term in duration." Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 4 (quoting U.S. Office of Management and Budget, Office of Federal Procurement Policy, "Publication of the Office of Federal Procurement Policy (OFPP) Policy Letter 11-01, Performance of Inherently Governmental and Critical Functions," 76 Fed. Reg. 56227, page 56236). A "closely associated" function is one that "may approach being" an inherently governmental function "because of the nature of the function and the risk that performance may impinge on Federal officials' performance of an inherently governmental function." Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 3-4 (quoting U.S. Office of Management and Budget, Office of Federal Procurement Policy, "Publication of the Office of Federal Procurement Policy (OFPP) Policy Letter 11-01, Performance of Inherently Governmental and Critical Functions," 76 Fed. Reg. 56227, page 56238).
[169] See Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 3.
[170] See U.S. Office of Management and Budget, Office of Federal Procurement Policy, "Publication of the Office of Federal Procurement Policy (OFPP) Policy Letter 11-01, Performance of Inherently Governmental and Critical Functions," 76 Fed. Reg. 56227.
[171] U.S. Office of Management and Budget, Office of Federal Procurement Policy, "Publication of the Office of Federal Procurement Policy (OFPP) Policy Letter 11-01, Performance of Inherently Governmental and Critical Functions," 76 Fed. Reg. 56227, page 56228.
[172] U.S. Office of Management and Budget, Office of Federal Procurement Policy, "Publication of the Office of Federal Procurement Policy (OFPP) Policy Letter 11-01, Performance of Inherently Governmental and Critical Functions," 76 Fed. Reg. 56227, page 56242.
[173] Federal Acquisition Regulation (FAR), "Subpart 7.5 - Inherently Governmental Functions," 7.503(c)(7).
[174] See Federal Acquisition Regulation (FAR), "Subpart 7.5 - Inherently Governmental Functions," 7.503(d)(5).

to agencies' own AI security teams (if any) so that contractors do not cross the line into directing federal employees.

In sum, government contracting is a complex topic, even outside the exigencies of an AI-related crisis, and thus more work will be needed on the surge utility of not only the mechanisms identified here but also others beyond the scope of this report—such as personal services contracts[175] and Defense Production Act (DPA) contracting authorities,[176] both of which come with their own considerations. Likewise, outstanding questions remain, perhaps most notably how to weigh the relative advantages and disadvantages of surging AI security talent using direct federal hiring versus contracting. To that end, potential future lines of inquiry include how to compare costs, which experts have long identified as a topic "beset by a host of challenges,"[177] as well as security issues and concerns that may stem from use of contractors, perhaps even including the physical security of AI talent.[178]

Outstanding Barriers to Surging AI Security Talent

There are multiple challenges that could frustrate efforts to surge AI security talent into government during an AI-related crisis. **Among the most urgent issues are security clearances, conflicts of interest restrictions, limited government hiring capacity and coordination, and divergent incentives between the government, private companies and other nongovernment entities, and AI security talent itself.** None of these challenges is intractable, but intra- and interagency planning and collaboration in advance of any crisis will be key. Thus, much work remains to be done, both on how to address the barriers discussed below, and on solving other challenges, such as how to structure budgets and funding to support a crisis-driven surge of AI security talent.

---

[175] The FAR specifies that a personal services contract (PSC) "is characterized by the employer-employee relationship it creates between the Government and the contractor's personnel." Federal Acquisition Regulation (FAR), "37.104 Personal services contracts," 37.104(a). PSC authority must be specifically granted to an agency by Congress. See Federal Acquisition Regulation (FAR), "37.104 Personal services contracts." PSCs are generally not considered to be federal employees for purposes of any law administered by OPM, though they may be treated as federal employees for other purposes. See, e.g., Agency for International Development Acquisition Regulation (AIDAR), "Appendix D to Chapter 7: Direct USAID Contracts With a U.S. Citizen or U.S. Resident Alien for Personal Services Abroad," 48 C.F.R. Chapter 7, subsection (4)(c)(1)-(2).
[176] See, e.g., 50 U.S.C. 4511 ("Priority in contracts and orders"), 4533 ("Other presidential action authorized," including "purchases of or commitments to purchase . . . a critical technology item, for Government use of resale"), 4558 ("Voluntary agreements and plans of action for preparedness programs and expansion of production capacity and supply").
[177] Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 24 (Appendix B).
[178] See Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 3; see also Frazier, "Is It Time for an AI Expert Protection Program?"

Security Clearances

The difficulty of obtaining security clearances—both in terms of **long timelines** and **strict standards**—has been a perennial challenge in federal hiring efforts.[179] Clearance-related barriers affect both individuals hired directly as federal employees and those retained as contractors; however, in the case of both "zero-hour" employees and contractors, the waiting time for security clearances can be paid upfront, before a crisis requiring them to begin work is imminent or has arrived.

Given the serious and potentially grave national-security implications of many of the crisis scenarios outlined in Section 1, AI security talent may need to receive Secret or Top Secret clearance to have access to the restricted information and systems needed to meaningfully contribute to response efforts. Moreover, ensuring that individuals are cleared to access information at those levels (even if there is not ultimately a "need to know" specific classified information designated as Secret or Top Secret during the particular crisis that presents)[180] is prudent because of the plausibility that such access will be needed, judging from the scenarios in Section 1, and the difficulty of predicting ahead of time the precise contours of the crisis that will emerge. The nature of the crisis-related intelligence also could require that talent have access to sensitive compartmented information (SCI), which is a subset of classified information that must be handled within formal access control systems, and thus necessitates additional review and approval before being granted.[181]

While average timelines for obtaining a Secret-level clearance run around several months, obtaining a Top Secret clearance with SCI generally takes a minimum of six months, though often closer to a year or more. Those long timelines are attributable to a variety of factors. First and foremost, security clearances are time- and resource-intensive, particularly with respect to the investigators necessary to conduct extensive interviews with the individuals seeking clearance as well as their family members, friends, current and former employers and colleagues, neighbors, and others. Wait times can also be an issue when already-cleared employees move between federal agencies. Specifically, a lack of interagency cooperation when it comes to security clearance "reciprocity," including differing interpretations of standards (e.g., on past drug use) and review processes between agencies, means that employees may have to wait weeks or months for their active clearance to be accepted by their

---

[179] As previously mentioned, this report does not address AI security talent who are not U.S. citizens, notwithstanding the fact that they comprise a significant portion of AI talent already studying or working in the United States. Nonetheless, it bears noting that another clearance-specific barrier to surging AI talent for crisis response is that non-U.S. citizens are not eligible for security clearances, though there can be other means of providing limited access. See Defense Counterintelligence and Security Agency, "Security Assurances for Personnel & Facilities: Limited Access Authorizations (LAA) for Non-U.S. Citizens."

[180] See Christensen, "Security Clearance Process: Answers to Frequently Asked Questions," 1.

[181] See, e.g., U.S. Department of Commerce, National Institute of Standards and Technology, "Glossary: sensitive compartmented information (SCI)"; U.S. Department of Commerce, Office of Security, "Access to Sensitive Compartmented Information (SCI)."

new agency. According to an expert consulted for this project, such delays in starting work can occur even when the clearance being transferred is at the same or a higher level than that needed for the employee's new position. Additionally, the typical cost to an agency for a new (i.e., first-time) Top Secret clearance currently starts above $5,000, though it can climb if there's a need for additional background investigation or priority processing.[182]

Contractors could face some additional challenges. While a contractor employee's security clearance generally remains in effect so long as that individual "remains continuously employed by the cleared contractor [entity] and can reasonably be expected to require access to classified information,"[183] a new clearance is needed when previous access to classified information has been inactive for more than 24 months.[184] Such new clearance might be granted comparatively quickly, given a prior favorable adjudication in the relatively recent past, though it could consume valuable time in a crisis. Nor do contractors represent any real cost savings to the government when it comes to clearances, as the government pays for clearances regardless of whether an individual will be a contractor or a direct federal employee. And while security clearance reciprocity between federal agencies is supposed to extend to contractor personnel as well,[185] one expert interviewed for this project said that, in practice, clearances of contractor employees can sometimes encounter even more challenges around portability between agencies.

"Interim" security clearances are a tool that can be used to shorten the amount of time that employees or contractors must wait before they can begin work. With an interim clearance, a prospective employee or contractor submits the requested paperwork, an initial review is conducted, and—if no concerns or complications arise—they may receive a provisional clearance allowing them to begin work while they await the final clearance decision. In terms of the wait time before beginning work, experts noted that interim clearances can shave several months off the timeline, though here as well, standards differ. Specifically, as one expert observed, some agencies make greater use of interim clearances than others, and **even when interim clearance is granted, they can translate into different access at different agencies, with some placing so many restrictions on the work that can be performed as to undermine its utility**.

Another major clearance-related barrier to surging AI security talent is the strict standards applied—at times inconsistently—by adjudicators. In particular, standards regarding drug use are an area where norms and expectations can differ to a degree, not only between the public and private sectors, but also between federal and state governments, as illustrated by the fact that most states have now legalized marijuana use in at least some

[182] See Defense Counterintelligence and Security Agency, "Billing Rates & Resources."
[183] Defense Counterintelligence and Security Agency, "FAQs–Facility Security Officers."
[184] Center for Development of Security Excellence, "Personnel Clearances in the National Industrial Security Program (NISP), Version 5," 10.
[185] Center for Development of Security Excellence, "Personnel Clearances in the National Industrial Security Program (NISP), Version 5," 11-12.

contexts,[186] and a handful have legalized certain psychedelics.[187] Furthermore, although the Office of the Director of National Intelligence (ODNI) has long issued generally applicable standards for evaluating risks in the context of security-clearance adjudications,[188] those guidelines—covering drug use and other issues, like foreign contacts—can nonetheless yield disparate results in practice. For example, **experts consulted for this project noted that some agencies show next to no tolerance for past drug use, while others use shorter "look-back" windows** that might discount drug use that occurred several years ago. While some experts referred to the possibility of "waivers" or similar tools whereby agencies might choose to grant a clearance despite past drug use based on a mission-critical need for a particular individual and/or other mitigating circumstances, it is yet another step that likely adds time to the clearance process. Likewise, to the extent that they are currently used, these sorts of waivers seem to operate in an ad hoc manner without consistent standards between or even within agencies.

Finally, **while beyond the scope of this paper, restrictions on granting security clearances to non-U.S. citizens also pose a barrier to clearing key talent**, as immigrants make up a large share of the top AI talent in the United States.[189] This is an area ripe for future research.

"Some clearance-related challenges could be at least partly addressed through planning ahead of an AI-related crisis."

In sum, clearance-related challenges are considerable, and some—such as standards around drug use—may require a change in mentality among agency officials.[190] But it bears noting that **other challenges, particularly those around long wait times, could be at least partly addressed by a degree of planning ahead of an AI-related crisis.** That is, by proactively recruiting and hiring (or contracting for) AI security talent, agencies could have them go through the clearance process, and, once cleared, keep them in intermittent or "zero-hour" status (and continuing to work in their regular, nongovernment jobs) until a crisis emerged. Then, agencies could call them to work immediately without needing to wait for the lengthy clearance process to conclude. While broader reform of the security clearance process is no doubt needed, this sort of advance planning remains integral to surging AI security talent in a crisis, whether via traditional federal hiring or government contracting.

---

[186] See Posard et al., "Updating Personnel Vetting and Security Clearance Guidelines for Future Generations," 19.
[187] See Wuthmann, "Massachusetts Voters Reject Proposal to Legalize Certain Psychedelic Drugs."
[188] See ODNI, "Security Executive Agent Directive 4 (SEAD-4)."
[189] See Oschinski et al., "Strengthening America's AI Workforce"; see also Defense Counterintelligence and Security Agency, "Security Assurances for Personnel & Facilities."
[190] See Posard et al., "Updating Personnel Vetting and Security Clearance Guidelines for Future Generations."

Conflicts of Interest

As explained above in the context of SGEs, federal conflicts of interest rules are significant, and the equity-related interests (e.g., RSUs) that are common in compensation packages at tech companies pose additional challenges. This is especially true in the context of surge hiring of AI security talent, where—given the short-term nature—individuals would likely be loath to divest any holdings for only brief periods of government service. Likewise, because of the short-term nature of a crisis-driven surge, this report envisions that AI security talent would maintain their nongovernment employment while temporarily serving in government. In that situation, conflicts of interest concerns are heightened because maintaining such ties arguably increases the risk that individuals could unduly influence the government to favor the services or products offered by those outside employers, or otherwise steer government matters to those outside companies' benefit.

SGE designation goes some way toward resolving conflicts-related challenges, though as discussed previously, **it is not a cure-all. As a supplement, agencies might consider designing "ethical walls" or "screens" and similar information-barrier protocols** to ensure that AI security talent, particularly those individuals coming from the private sector, does not have access to certain information or decisionmaking processes that could create actual or apparent conflicts of interest. Waivers are another tool that agencies might be able to use to address some conflicts, though they are likely to be fact-specific and considered only on a case-by-case basis. Establishing such processes undoubtedly requires forethought and careful pre-planning by agencies' HR teams and ethics attorneys, to ensure that conflicts are resolved and mitigated to the extent possible ahead of any crisis, so as to not lose valuable time before surged talent can begin work.

As referenced in the discussion of SGEs above, **contractor personnel generally are not subject to the conflicts of interest restrictions applicable to direct federal employees.** In particular, a contractor employee's equity in the contracting company that they work for does not pose the same concern that equity-related interests do in the joint government and private-sector employment scenario discussed above. **That said, "organizational" conflicts of interest remain a live concern** when it comes to contractors, and so agencies still must take steps to make sure that, e.g., contracting companies do not receive an unfair competitive advantage in gaining future contracts through their existing arrangements with the government, or through efforts of their contractor personnel that might influence the procurement process.[191]

Government Hiring Capacity and Coordination

Akin to the way in which the government's shortage of personnel and other resources can frustrate efforts to expedite security clearances, the government likely lacks the full extent of

[191] See Federal Acquisition Regulation (FAR), "Subpart 9.5 - Organizational and Consultant Conflicts of Interest."

capacity and expertise needed to conduct large-scale recruitment, surge hiring, and onboarding of AI security talent. Thus, when discussing how to surge AI security talent into the government for crisis response, **it is vital to also consider the role of "AI-enabling" talent ahead of any crisis.** Such "enabling" talent includes hiring managers, contracting officers, and others involved with obtaining and training new personnel and ensuring that they are set up for work (e.g., IT teams).

As one expert noted, the government has historically outsourced a sizable portion of technology-related work, such as software development. Moreover, as a relatively new field, AI—and the relevant qualifications and experience for such roles—may be even less familiar to federal hiring teams than other tech disciplines. And to the extent that not all highly skilled AI researchers necessarily have advanced degrees, hiring teams may lack another measure commonly used in making federal hiring decisions.[192] As a consequence, **hiring managers and other employees responsible for agencies' hiring efforts may not currently have the AI-fluency needed to recruit and evaluate AI security talent quickly and at scale.** On top of that, experts described federal hiring teams as frequently understaffed even during normal operations, further reducing bandwidth to surge AI security talent before or during a crisis.[193] A similar dynamic likely exists in the contracting context as well, where personnel tasked with procuring and managing contracts for AI security expertise might themselves lack the specialized background, time, or tools necessary to fully assess performance under contracts and to hold contractors accountable when needed.[194]

Beyond recruitment and deciding whom to hire, challenges persist at the onboarding and training stages. Onboarding imposes numerous time burdens on an agency, including in the form of HR functions like facilitating building access and badging. There are also various IT functions, such as setting up access to digital systems (including file repositories and research databases), and issuing hardware like computers and cell phones, all of which normally take a few weeks. And while training on topics like cybersecurity, ethics, and safeguarding classified information are normal parts of entering government service, it is possible that agencies would need to devise and implement additional, more hands-on training to effectively surge AI security talent. Such training could be key to ensuring that AI security talent—many of whom may not have prior experience working inside the federal government, and who may not have the luxury of getting acclimated over the course of several weeks—can hit the ground running in a crisis. Building out that sort of training program, so that surged talent is fully aware of the chain of command and other important aspects of government operations, would require its own timeline and resources.

---

[192] See The White House, "AI Talent Report."
[193] See, e.g., Partnership for Public Service, "Roadmap for Renewing Our Federal Government."
[194] See Halchin, "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues," 20 (noting, in the context of the IC, the "challenge" of "maintain[ing] an acquisition workforce that has the capacity, in terms of size, expertise, and training, to oversee the many contractors and contractor employees performing work for the intelligence community").

Coordination-related challenges also abound. Government hiring is typically decentralized, handled on an agency-by-agency basis. Without real-time information sharing, agencies could end up competing for the same limited pools of talent. While centralizing surge hiring in a single agency might yield efficiencies of scale, since personnel at that agency would gain repeated experience with the necessary hiring and onboarding actions, it might create a bottleneck by overburdening those teams, and could make it more difficult to ensure that AI security talent is dispersed as needed across various affected agencies when a crisis emerges. Relatedly, centralizing AI-related surge-hiring efforts in an office within the White House presents its own challenges, as such offices typically lack the hiring (and contracting) expertise, bandwidth, and funding that comparatively larger agency teams possess.

Incentives

Finally, yet another barrier to surging AI security talent is that incentives for the relevant actors—namely, the government, private companies and other nongovernment entities, and individuals—tend to push in different directions, frustrating efforts to balance costs and benefits and to encourage broad participation.

In general, a federal agency's incentives are to minimize the cost—both in time and money—in obtaining AI security talent, while not sacrificing the quality of such talent. The government also has an interest in maintaining oversight and control over the work done by AI security talent, to ensure that it is directed and coordinated in the ways that the government needs, in service of the public good.

> *"Incentives for the relevant actors—government, private companies and other nongovernment entities, and individuals—tend to push in different directions."*

Companies and other nongovernment entities have numerous potential incentives. They are incentivized to **reduce the disruption to their normal activities** that comes from their employees' time spent in government service, while also reaping the benefits of their employees' experiences and contacts inside the government. They are also interested in **maintaining good government relations**, which could lead to greater access, information, and trust, in turn benefiting the company or organization. Of course, **private companies will have a profit motive**, which could manifest in a variety of ways, from a desire to obtain lucrative government contracts, to an aversion to entering into such arrangements given high compliance costs and potential reporting or oversight requirements. Companies may also have the incentive to protect their intellectual property. Moreover, companies—particularly if their models were somehow implicated in an AI-related crisis—would likely derive value from being part of the government's response, which could mitigate the fallout from what might otherwise be a serious public-relations problem. Relatedly, avoiding liability for AI-related crises might be another potential incentive for companies and other entities. **Public benefit corporations, academic institutions, and civil society might also be particularly motivated to serve the public interest during an AI-related crisis.**

At the level of the individual, incentives include the **additional prestige** or, for some, extra income that may come with government service. Furthermore, a stint in government—particularly during a high-stakes crisis situation—could provide individuals with **valuable experience and insight** into how the government operates in making decisions and coordinating across agencies. Such experience could also expand individuals' professional network, increasing their contacts within both government and the other categories of AI security talent that would be surged into government during a crisis response. Last but not least, individuals might be motivated by **a sense of purpose to serve in government for short periods during an AI-related crisis**, particularly given the grave national-security implications of the crisis scenarios outlined in this report.

Finding equilibrium between these conflicting incentives poses a challenge, despite consensus that the government's response to an AI-related crisis would be enhanced by access to AI security talent from outside of government. Specifically, where the government seeks to reduce costs, companies and individuals may lack the financial incentive to participate in the project of surge hiring. What's more, financial incentives of any sort may not be enough to attract the employees of frontier companies, given the vast pay differential between the companies and government. While individuals might be drawn to short-term government service by the experience and contacts it offers, companies and other nongovernment employers could be loath to relinquish their talent for even short periods of time. Moreover, the government might differ with companies and other nongovernment entities over who primarily directs the activities of employees, how much time is dedicated to public service, and in what context competitors (or their employees) have any access to rival IP.

4 | Policy Solutions for Surging Talent

The U.S. federal government has several policy options to address critical gaps in AI security talent during crisis situations. The policy suggestions outlined below propose:

1. **Conducting agency inventories of existing in-house AI security expertise**,
2. **Establishing a national AI reserve corps**,
3. **Leveraging contract-based solutions**,
4. **Establishing a new hiring authority**,
5. **Reforming the Intergovernmental Personnel Act**,
6. **Updating the security clearance process**, and
7. **Building the capacity to hire**.

These recommendations can be adopted individually or in combination, with the White House or OPM providing coordination and delegating specific responsibilities to agencies. Given the depth of the talent gaps, however, any single one of these recommendations—with the possible exception of establishing a National AI Reserve Corps—is likely insufficient for overcoming current barriers to rapidly onboarding AI security talent. As such, we recommend integrated reform across hiring and clearance systems to produce the surge capacity necessary for major AI security emergencies.

“While comprehensive reform takes time, several immediate actions could accelerate the government’s ability to access talent during crises without requiring legislative changes or lengthy bureaucratic processes.”

While comprehensive reform takes time, several immediate actions could accelerate the government’s ability to access talent during crises without requiring legislative changes or lengthy bureaucratic processes. These include conducting systematic inventories of existing AI security expertise across agencies, expanding the use of contract vehicles to allow for the redirection of relevant contractors for crisis response, and OPM creating a new special hiring authority within Schedule A. Although these steps are not complete solutions in and of themselves, they could be implemented within months at the agency level and could provide meaningful improvement in surge capacity while longer-term reforms advance.

Conducting Agency Inventories of Existing In-House AI Security Expertise

As a first step toward building surge capacity, federal agencies should conduct systematic inventories of their existing in-house AI security expertise and report their findings to the Office of Personnel Management. A clear understanding of current staffing levels, technical competencies, and mission-specific gaps will allow surge mechanisms to complement, rather than duplicate, existing federal capabilities. These inventories should extend beyond headcounts to include details on staff clearances, relevant training, and cross-domain expertise (e.g., AIxcyber, AIxbio). Regularly updated inventories would also strengthen interagency coordination by giving the federal government visibility into where expertise resides across agencies, enabling more efficient mobilization when crises arise.

As part of this effort, agencies should build on federal initiatives to categorize AI talent currently serving in government, and to forecast future talent requirements.[195] **Any comprehensive inventory of the federal AI workforce should explicitly include AI security crisis-response roles** to ensure the government can rapidly identify and deploy specialized expertise when emergencies demand immediate action.[196] A shared lexicon can ensure federal-level inventories of AI security talent maintain consistency across interagency, improving interoperability and mobilization capabilities. The Typology of AI Security Talent (see Appendix A) provides a starting point for creating that shared lexicon.

Establishing a National AI Reserve Corps

A National AI Reserve Corps would create a pre-vetted pool of cleared technical experts from industry, FFRDCs, academia, nonprofit research organizations, and independent researchers who can be rapidly activated during AI-related crises. The corps would operate through a dedicated program office within an executive agency and would recruit specialists based on their skill sets, preemptively provide security clearances to enable swift deployment, and conduct regular joint training exercises to maintain operability. Reservists would be designated as Special Government Employees (SGEs) for short-term stints and would receive compensation and employment protections similar to those given to military reservists.

---

[195] For example, the AI in Government Act of 2020 required OPM to estimate the current number of federal employees working in positions related to AI and categorize AI competencies that exist in government or may be needed for future AI-related positions. OPM subsequently released three publications to provide guidance on classifying and hiring AI talent (NSF, "Artificial Intelligence Scholarship for Service Initiative," 13). Furthermore, as of this writing, there is pending legislation to develop a framework for AI talent that is similar to National Initiative for Cybersecurity Education (NICE). See Peters, "Senator Peters Reintroduces Bipartisan Bill to Boost America's AI Workforce Pipeline."

[196] While the National Initiative for Cybersecurity Education framework does include a competency for AI Security, a more detailed categorization such as the Typology of AI Security Talent (see Appendix A) would help the federal government more rapidly and accurately identify the specific competencies needed to respond to a given crisis. See NICCS, "NICE Framework Competency Areas."

While establishing a National AI Reserve Corps requires upfront investment, and maintaining and deploying it requires significant interagency coordination, such a model could offer advantages over other hiring measures. The reserve-corps structure provides the government real-time flexibility in surging this talent into and out of crisis response teams as incidents evolve, while allowing streamlined updates to the reserve's composition as technical AI security competencies advance. Likewise, it provides a mechanism through which agencies can better coordinate access to and deployment of talent across the government, rather than competing with each other for the same limited talent pools, as might happen if resorting to normal ad-hoc hiring practices. The government also benefits from specialists whose skills remain cutting-edge through private sector work (as with other surge hiring using intermittent schedules), without bearing the costs of continuous training and development. The reserve corps structure would also allow the federal government to coordinate joint drills with reservists and interagency teams to maintain operational readiness, a benefit that is harder to obtain when using other surge mechanisms, such as contract-based solutions. Overall, the National AI Reserve Corps could provide the government with immediate access to the nation's top AI security talent without many of the coordination failures that undermine other civilian surge mechanisms.

Structure: The U.S. government should consider which federal body is best positioned to operate a National AI Reserve Corps. Several candidates stand out, though each comes with advantages and tradeoffs:

- *White House*: There is precedent for housing technical talent initiatives at the White House (e.g., the U.S. Digital Service). White House affiliation also carries significant prestige that could help attract high-caliber AI security experts. However, the White House is primarily a policy and coordination body, and lacks the deep operational infrastructure of federal agencies. In addition, programs located in the White House can be especially vulnerable to shifts in administration priorities, raising concerns about long-term continuity.
- *Department of Energy*: DOE hosts one of the government's largest concentrations of scientific and technical expertise and has longstanding experience partnering with academia and industry through FFRDCs and cooperative agreements, which could facilitate integration of outside experts into a reserve structure. DOE's role in safeguarding critical infrastructure (e.g., the energy grid) also overlaps with likely AI crisis domains. However, its primary mission is not emergency incident response.
- *Department of Homeland Security*: DHS holds statutory authorities for emergency preparedness, critical infrastructure protection, and cybersecurity, aligning directly with likely AI crisis domains. It also has experience with surge personnel models (e.g., the Surge Capacity Force) and maintains operational relationships with relevant private sector stakeholders via CISA and the Joint Cyber Defense Collaborative. However, DHS

resources are currently focused on other priorities, and DHS's reputation may hinder recruitment from some key talent pools.[197]

- *National Institute of Standards and Technology*: NIST hosts CAISI, a small but agile body of AI expertise that both understands the nature of the potential crises at hand and the talent that might need to be surged to respond to them. However, CAISI's limited resources, already heavy mandate—particularly following the AI Action Plan—and relative lack of experience in crisis response raise concerns about whether it has the capacity to take on a new responsibility of this scale.

A dedicated program office within one of the above federal bodies—perhaps led by a National AI Incident Coordinator—would administer the National AI Reserve Corps. Core functions of the office would include: recruiting reservists who possess priority skills and competencies (see Appendix A); training; managing security clearance processes and ethics compliance; and coordinating interagency activation requests and deployments in consultation with partners such as OSTP, the NSC, and other federal crisis management entities.

Compensation: A reserve corps will need to balance two core objectives: offering compensation that is competitive enough to attract highly skilled AI experts, and structuring pay in ways that remain economically feasible for the federal government. During activation, one option would be to compensate reservists at GS-equivalent rates, consistent with other federal surge programs. Alternatively, Congress could authorize a capped stipend or allow higher-than-GS pay bands under a new special hiring authority, creating flexibility to draw from talent pools where opportunity costs are steep. Outside activation, reservists would continue working for their home institutions, which would help reduce government costs. In cases where federal pay lags behind private-sector salaries, policymakers should strongly encourage employers to supplement federal compensation to the maximum extent legally allowed to make employees whole, including through continued equity-based compensation, thereby making participation more viable.[198]

Clearances: Even when expedited, security clearance processing can take months, potentially standing in the way of rapid incident response. To avoid these delays, security clearance sponsorship and processing could occur preemptively, as soon as hiring (or contracting) decisions are made and finalized, with support from lead and/or participating agencies' personnel security infrastructure.

[197] A number of researchers at frontier AI companies interviewed for this project specifically stated that they, and likely their peers, would be less likely to work for a DHS sponsored program, citing DHS's reputation for bureaucracy and objections to the work it performs.

[198] As previously explained in Section 3, it is generally only possible for federal employees to receive compensation from outside entities for government work with SGE designation (or if they are unpaid experts/consultants, or IPA detailees whose salaries continue to be paid by their regular outside employers). For this reason and others, SGE designation is another part of the reserve corps recommendation.

Training: A number of former U.S. government officials interviewed for this report emphasized that civilian reserve corps are particularly vulnerable to *coordination failure*. Civilian experts are drawn from diverse institutions, often with little shared experience operating with each other or within federal crisis frameworks. Without regular joint training, reservists risk confusion over lines of authority, mismatched expectations about roles, and slow integration into government structures. To mitigate this, the coordinating office should consider organizing regular—perhaps biannual—training exercises for reservists, in which reservists and relevant federal agencies would participate in advanced technical simulations of crisis scenarios and in post-exercise debriefings with interagency representatives to iterate on existing incident response frameworks. Any such trainings would need to be designed in consultation with the relevant federal agencies and potential reservists to ensure that trainings meet all stakeholders' expectations of quality, value, and feasibility. Trainings could be held on the margins of major AI conferences, where many reservists would already be in attendance, to reduce potential barriers to participation.

SGE Designation: Reservists would be classified as Special Government Employees (SGEs), allowing them to serve up to 130 days per year while subject to somewhat less restrictive federal conflicts of interest restrictions (see Section 3).[199] This designation would enable greater participation from experts—in particular, those employed by the private sector, but also academic bodies and nonprofits—while safeguarding against real or perceived conflicts of interest.

Authorities and Appropriations: Once a lead agency is identified, Congress could enact authorizing legislation establishing a Program Office within that agency and instructing it to recruit and maintain a National AI Reserve Corps. Effective legislation would also empower the Program Office to appoint reservists under a special hiring authority (see "A New Hiring Mechanism for AI Security Talent") and define protocols for activation in response to crises or interagency requests. This authority could be modeled on existing frameworks, such as the DHS Capacity Force.

Alternatively, the Executive Branch could take steps to establish the National AI Reserve Corps, with OPM using Schedule A to create a reserve corps authority, and a lead agency standing up a Program Office subject to any required Congressional consultation or notification requirements. While Executive Branch action to establish a reserve corps might be comparatively faster, legislative action by Congress remains the preferred route for the

[199] Though this report is focused only on short-term surges of AI security talent into government (likely on the order of days or weeks), it's worth noting that, per OGE, if an employee "unexpectedly serves more than 130 days during the ensuing 365-day period, the individual still will be deemed an SGE for the remainder of that period," although the agency must reevaluate whether SGE designation continues to be appropriate for the next 365-day period. Office of Government Ethics, "Memorandum dated February 15, 2000, from Stephen D. Potts, Director, to Designated Agency Ethics Officials, General Counsels and Inspectors General Regarding Summary of Ethical Requirements Applicable to Special Government Employees," 5. Thus, SGE designation is unlikely to hinder response efforts to even a more protracted AI-driven crisis than those contemplated here, though it could affect whether SGE designation would continue to be a feasible option going forward.

additional flexibilities that Congress (and only Congress) might confer, such as flexibility on pay to boost recruitment efforts, which could be included as part of a new legislatively granted hiring authority. Likewise, given that Congressional buy-in on establishing a Program Office is likely necessary in either scenario, authorizing legislation could come with the advantage of increased access to the OTA authority discussed in Section 3, which NIST and White House offices like OSTP do not currently have.

To ensure that there is adequate funding for reservists' compensation and training programs, Congress should also provide dedicated appropriations to fund the National AI Reserve Corps. Funding could be further supplemented through interagency agreements in which partner agencies reimburse the lead agency for reservist support. During declared emergencies, appropriations from contingency or disaster response funds may also be available to address additional, evolving needs.

Incentives for Participation: A central challenge for any AI Reserve Corps is that many experts may have little incentive to join, given significant pay disparities, time constraints, and potential conflicts with their primary employers. Policymakers should therefore consider what factors could make participation more attractive. Based on our interviews, a few points of emphasis come to mind. For individuals, possible levers include emphasizing the opportunity to contribute to high-stakes national missions, developing government experience and cross-sector credibility, and building professional networks. For employers, participation might be framed as a demonstration of public-interest commitment, a way to strengthen relationships with government, or an avenue for select staff to gain security clearances and crisis-response experience. Additional measures such as modest financial support, tax incentives, or limited non-competitive eligibility for later federal roles could also be explored, along with recognition mechanisms for employers that support employee participation.

Making Greater Use of Contract-Based Surge Capacity

Another approach to surging AI security talent is for the federal government to make greater use of contracting with the companies and organizations that employ such professionals (see Sections 2 & 3). **The government can do this by identifying AI security talent included in current federal contracts, adding crisis-response clauses to future contracts that will include relevant personnel, and creating new contract structures explicitly designed to address AI security surge needs**. Particularly with respect to new contracts, agencies should move proactively to ensure that agreements are in place ahead of a crisis, given that lead times on establishing new contracts are often several months or more.

The federal government may be able to draw from existing contracts with external entities to field ready-to-surge teams of AI security talent. First, agencies should map AI security talent currently working for the government via federal contracts. Once talent is available to the government through contractual arrangements, U.S. government contracting experts advise that it can be easier to redirect these contractors to different lines of effort, including interagency task forces responding to AI security crises, as compared to executing new contracts. These individuals are already cleared, integrated into government systems, and paid for through existing contracts. That said, the terms of these preexisting contracts must allow for work to be redirected based on such exigencies. While one interviewee noted that such terms are common, it is important for agencies to map these agreements ahead of time both to understand which areas of competency might be covered (and where there are gaps) and to verify whether redirecting such personnel is possible within contractual terms.

To institutionalize this approach going forward, agencies should incorporate crisis-response language when possible into future contract statements of work for agreements that include AI security talent. When awarding new contracts that include AI security talent, agencies can include provisions allowing for personnel with AI security expertise to be redirected into crisis response activities when the need arises. This approach ensures that talent acquisition occurs during normal operations and not in a high-pressure crisis situation. However, on its own, relying on contracts entered into for normal operations and services will likely not be sufficient to provide the comprehensive range of talent necessary for responding to the full array of possible crisis situations. Furthermore, some types of arrangements—such as the OTAs discussed in Section 3—may not be susceptible to being redirected in this way, given statutory and regulatory limits on how they can be used.

The government can further enhance its surge capacity through "zero-hour" contractor arrangements, in which AI security experts are placed on retainer and bill only for hours served while remaining available for rapid activation during a crisis. As with "zero-hour," intermittent-schedule federal employees (see Section 3), this approach allows talent to be cleared and badged in advance, reducing delays associated with clearances and administrative onboarding. There is already precedent for such an arrangement—FEMA awards "advance contracts for goods and services" that are competed and awarded before any disaster occurs, but which allow for rapid deployment when crises occur.[200]

Another possible approach is to use an indefinite delivery, indefinite quantity (IDIQ) contract structure with multiple prime contractors who then compete for task orders to provide the government with services from their cleared AI security talent. Under this arrangement, the government could award IDIQ contracts to multiple companies that employ AI security talent, and so begin the process of clearing necessary talent. When surge capacity is needed, agencies could rapidly issue task orders to these companies based on terms established during the initial contract award, eliminating the need for time-consuming execution of an initial contract during crisis response. This approach allows the government to activate multiple contractors simultaneously if needed, scaling response capacity based on the severity of the crisis. However, establishing IDIQs before a crisis is critical, given that initial IDIQ contracts generally take months to put into place. Agencies' contracting teams should also take care to ensure that IDIQs provide for adequate competition for task orders, which could prove

[200] FEMA, "Advance Contracts of Goods and Services."

challenging in areas of competency that are dominated by only a few companies. Such competition is important to avoid vendor lock-in and maintain downward pressure on pricing.

In general, contracting approaches could be a practical way to surge AI security talent in short-term crisis scenarios in which traditional hiring mechanisms could be too cumbersome, or where those mechanisms (along with their attendant ethics requirements) might complicate full access to certain talent pools, such as the private sector. However, as discussed in Section 3, the crisis-response efforts of AI security talent could qualify as "critical" functions or "closely associated" with inherently governmental activities, meaning that agencies must ensure that they have the capacity, both in terms of oversight processes and their own in-house expertise, to appropriately manage contractors' activities. That regulatory requirement provides yet another reason why advance planning is necessary to ensure contractors can effectively respond to crises without assuming core governmental functions, such as directing or controlling federal employees or being treated as de facto decisionmakers by agency officials.

Despite these limitations, a contract-based approach offers some advantages for surging talent. Whereas hiring efforts might require reform (potentially including legislative action) on issues like ethics requirements in order to surge AI security talent quickly and at scale, contracting mechanisms avoid some of these barriers and so might be more fit for purpose without the need for as extensive intervention (though, notably, they face many of the same challenges around issues like security clearances). Furthermore, the contracting option aligns incentives across the government, industry, and individual professional level. Private sector companies earn revenue and strengthen government relations, and their employees gain public sector experience while maintaining their private sector employment and compensation. The government builds out its AI security surge capacity quickly without necessarily requiring legislative intervention or the often lengthier processes associated with traditional hiring, as long as it has laid the groundwork for activating contracts. This alignment of interests makes contracting a practical approach for addressing AI security talent gaps in crisis situations.

Expanding Hiring Authorities to Access Private Sector Talent

A New Hiring Mechanism for AI Security Talent

Whether as part of an AI reserve corps or as a standalone measure, Congress should establish a new special hiring authority dedicated to hiring AI security talent for crisis response. Such an authority would have several advantages over existing authorities.

- **First, a new special hiring authority promotes efficiency and coordination.** A special hiring authority would likely increase the speed of the hiring process before—and even more significantly, during—a crisis, because it would make it obvious to agency hiring teams which authority to use, rather than requiring them to compare and contrast the limitations of existing authorities to determine which is the best (or least worst) option. While hiring managers may sometimes be hesitant to use new authorities in lieu of familiar ones, that is unlikely to be the case here, where the authority is narrowly tailored to a specific purpose and need.

- **Second, it facilitates access to all potential pools of AI security talent, particularly in the private sector.** With access to a single authority, rather than resorting to the patchwork of existing ones, agencies could embark on streamlined hiring efforts without potentially excluding the unique expertise of the private sector (as IPAs currently do) or wading into potentially limiting (and time-intensive) line-drawing questions regarding what work an agency's regular employees do or do not perform (as required with the expert/consultant authority).
- **Third, it encourages agencies to take a proactive approach to planning for an AI-related crisis.** A special hiring authority would likely prompt action by signaling to agencies that hiring AI security talent for crisis response is a Congressional priority, nudging agencies to do the time-consuming work of recruiting, selecting, clearing, and onboarding AI security talent (potentially on intermittent "zero-hour" status) before a crisis actually occurs.

Taken together, a new special hiring authority would notably enhance both the effectiveness and efficiency of surge hiring efforts to respond to an AI-related crisis.

Congress routinely creates special hiring authorities in authorizing legislation, most notably the annual National Defense Authorization Act (NDAA). These grants of authority are generally drafted in straightforward language, and often quite brief—frequently only a sentence or two.[201] Crucially, special hiring authorities typically specify that they are excepted from the default laws governing competitive-service hiring, which require time-consuming hiring procedures.[202] They often also include an exemption from the Title 5 provisions governing pay, granting discretion to the agency to set the rate of pay. Combined, these exclusions would enable agencies to effectively and nimbly identify top talent and make offers that are sufficiently attractive to secure it. If Congress so chose, it could extend this special hiring authority to all agencies, or identify a specific agency or set of agencies to exercise the authority.[203]

---

[201] See, e.g., 22 U.S.C. 2651a(i)(5) (authorizing the Secretary of State to "appoint up to 25 employees to cyber positions" in a Congressionally created Bureau of Cyberspace and Digital Policy). This authority was granted in the NDAA for FY2023 (Public Law 117-263, Dec. 23, 2022, section 9502).

[202] In addition to specifying that the Title 5 rules regarding "appointment" (i.e., the ordinary competitive-service hiring procedures) do not apply, any special hiring authority for AI security talent should also specify that the Title 5 provisions regarding "termination" do not apply. This second exception is important to ensure that the government can move quickly to remove any individuals who fail to meet the obligations of government service.

[203] If deemed necessary, Congress also has the option to cabin agencies' use of the authority in other ways, such as by specifying a ceiling on the number of positions or including a sunset clause, i.e., an expiration date (generally several years in the future) on which the authority expires unless extended or reauthorized by Congress. While sunset clauses can complicate agencies' planning efforts, they do give Congress a clear opportunity to revisit an authority periodically, at which point it can consider not only whether to extend the authority but also whether to expand its scope based on past lessons and any changes in circumstances. For example, Congress could decide to raise any limits on the number of positions, or to extend the authority to additional agencies.

That said, a special hiring authority alone cannot address all the barriers to surging AI security talent during a crisis. Hiring qualified talent at scale requires funding, and while agencies can likely repurpose some existing discretionary funds toward this end, it would be wise for Congress to appropriate funding expressly for this purpose. Hiring also requires coordination and planning to identify, recruit, clear, and onboard talent. Doing so quickly and effectively will also require increasing the resources devoted to crucial AI-enabling talent like hiring managers (especially ones with expertise in the markets and skill sets identified in the Typology of Talent). In both instances, Congress has a role to play in both funding and guiding the way in which the Executive Branch builds this capacity.

Alternatively, OPM has some ability to establish new hiring paths through Schedule A (as well as Schedule B and Direct-Hire Authority; see Appendix B). It can do so on either a governmentwide or agency-specific basis. For example, OPM could authorize all or some agencies to use Schedule A to hire AI security talent for crisis response. OPM authorization is an appealing route, given that it likely could be granted more quickly than a legislative fix. That said, OPM generally lacks the freedom that Congress has to craft authorities that provide additional flexibility on matters like pay, which are largely set by statute. Thus, while a Congressionally granted special hiring authority is likely still the preferred option, OPM authorization would be another viable—and welcome—route.

Intergovernmental Personnel Act Reform

A new hiring authority specific to AI security talent would be optimal to facilitate surge hiring of experts across industry, academia, and civil society. However, the U.S. government may consider adapting existing authorities as an alternative to creating a new hiring authority. Of existing options, the Intergovernmental Personnel Act (IPA) is among the best-suited, given its established precedent for bringing technical talent into government, possibility for cost-sharing agreements between sending organizations and the government, and compatibility with intermittent or short-term appointments.[204] However, the IPA authority, as currently written, has narrow eligibility requirements that would hamper its efficacy for a surge-hiring use case, as detailed in Section 3. In particular, **for-profit organizations—where AI security talent is disproportionately located—are not IPA-eligible under current federal law.** Therefore, the IPA program would require reform to be a fully viable authority for surging AI security talent in a time of crisis.

Congress should consider expanding IPA eligibility to include private-sector, for-profit entities that employ AI security talent. While AI security expertise exists across industry, academia, and civil society, much of the talent is concentrated in industry (see Section 2). Federal agencies have previously recognized the need to ease barriers to rotating private-sector

[204] GAO, "Improved Guidance Could Help Federal Agencies Address Skills Gaps and Maximize Other Benefits," 7.

expertise into government,[205] and in the National Defense Authorization Act for Fiscal Year 2017, Congress approved a separate IPA-like authority allowing for-profit private-sector employees to serve short-term stints at DOD—called the Public-Private Talent Exchange (PPTE).[206] Several other intra- and extra-governmental organizations have called for IPA expansion to include eligibility of for-profit organizations across not only the DOD, but also federal civilian agencies.[207]

A significant barrier to IPA expansion is legitimate concern over corporate influence and agency capture.[208] In a general sense, these concerns apply here as well. Crisis response will involve access to critical information and direct collaboration with government actors who may be involved in the regulation of private companies. Moreover, at least some crisis scenarios will implicate the financial interests of individuals employed by these private companies.

That said, several factors unique to surge capacity limit many of the concerns presented by IPA expansion in other contexts. Most notably, surge capacity is naturally limited in scope and primarily involves tasks that are less prone to abuse. As noted in Section 1, most of the needs envisioned by this report involve performing technical tasks and solving discrete problems, rather than making policy or regulatory judgments. These tasks present fewer worries of influence than more "general-purpose" IPA arrangements. Additionally, surge capacity is limited in duration and ad hoc, reducing the likelihood of longer-term influence or repeat self-dealing.

Congress will still likely want to limit the risk of undue influence, and could do so by imposing additional time limits on assignments, requiring that agencies disclose use of the authority to Congress to promote oversight, or expressly prohibiting hires made under this authority from performing certain tasks (e.g., any tasks related to discretionary policy determinations or rulemaking). Similar restrictions can be found in the PPTE and related regulations and may serve as inspiration for these limitations.[209]

Congress may also want to make these changes without altering general use of the IPA authority. Drawing on the example of the PPTE, Congress could develop a public-private talent exchange program distinct from the IPA program specifically for undersupplied AI security talent. The key difference would be the provision's scope: while the PPTE is limited to a single agency (DOD) by statute, and otherwise does not restrict when the agency can use the

---

[205] Boyd, "DoD 'Force of the Future' includes civilian Digital Service team."
[206] DCPAS, "Public Private Talent Exchange Program."
[207] MITRE, "Symposium on the Federal Workforce for the 21st Century," 18; Partnership For Public Service, "Trading Places: The Benefits, Challenges and Potential of Federal Public-Private Talent Exchanges."
[208] See the following for an example of prior concerns about corporate influence through an IPA program: Grassley, "Grassley Launches Sweeping Review of Program Allowing Privately Employed Individuals to Serve in Federal Government Roles."
[209] DCPAS, "Public Private Talent Exchange Program."

authority, an expansion of the IPA to target AI security talent would likely need some limiting principle. Promising limiting conditions include:

- **AI-relevant expertise**: The individuals hired possess the "knowledge, skills, and abilities to be considered subject matter experts"[210] in occupational fields relevant to AI security. Congress would need to define the term "AI security," or a related phrase, in statute or otherwise instruct agencies on the scope of this power.[211] This would likely be done with input from relevant agencies and stakeholders in order to identify the specific talent profiles Congress wishes to preference. The Typology of Talent included in Appendix A may be a useful starting point.
- **Unmet need**: The agency cannot obtain the required expertise or an adequate volume of that expertise through other available hiring authorities. This would include inability to hire adequate talent due to an acute, urgent need.

Congress could further narrow such an authority by requiring that hiring is necessary, appropriate, or related to addressing issues of:

- **National security, public health, or public safety**, or
- **AI-related emergencies** as defined by statute.

These additional criteria have both benefits and limitations. By narrowing the authority to more severe situations of public need, Congress can prevent undesired expansion or abuse of this authority. However, as history suggests, tying authorities to emergency circumstances can encourage agencies to expand the meaning of those terms and invoke emergency powers in more contested circumstances. Requiring the existence of an ongoing emergency would also prevent the sort of proactive recruitment highlighted in this report. For these reasons, Congress may wish to narrow this authority via other levers (i.e., length of assignment, scope of tasks, etc.) or require only that the *tasks* performed by hires relate to issues of national security rather than tying the authority to the *existence* of an emergency at the time of hiring.

By adopting some set of limiting criteria, Congress can help agencies access critical AI security talent concentrated in private industry without straying too far from the IPA's more general exclusion on obtaining private-sector talent.

To operationalize this expanded authority, the government would proactively identify specific companies that could serve as sources for AI security talent in advance of any crisis, begin the process of selecting their interested and qualified employees, and start executing the necessary agreements for intermittent work, activating

[210] See DCPAS, "Public Private Talent Exchange Program."
[211] Congress has several options in defining a phrase like "AI security" talent. It can enumerate specific categories of covered skills, set out more general criteria, or instruct agencies to issue regulations defining what skill sets or roles are in scope. In any event, the key task will be to balance flexibility (to ensure that a definition does not become obsolete or fail to update to new circumstances) while providing sufficient direction to avoid both abuse and legal challenge.

assignees only on an as-needed basis. As detailed in Section 3, organizations cannot be preemptively certified as IPA-eligible (i.e., before they have an individual going through the IPA assignment process), though once certified as eligible, that certification is permanent and may apply across the federal government, meaning other agencies need not separately certify.[212] Thus, the government would be wise to preemptively engage with organizations (including entities that have not previously gone through the IPA certification process, such as newly founded nonprofits) to clarify their potential role in surge response and identify employees who could be mobilized. By initiating eligibility certification and executing at least some IPA assignment agreements in advance, the government could reduce delays in establishing any additional assignments with those same entities during a crisis, saving valuable time.

To make this approach effective, Congress may also need to adjust existing ethics rules for the unique context of surging key private-sector AI security talent for short-term government service. Even if for-profit entities are made IPA-eligible, the current ethics requirements applicable to IPA assignees—and federal employees in general—could disincentivize private-sector employees, including AI security talent from the nation's frontier companies, from participating in an IPA-enabled program. Ethics requirements would need to account for the tech industry's unique compensation plans, including equity-based pay with vesting schedules, and post-duty restrictions on the types of projects a returned IPA assignee (or private-sector employee serving in government under any other authority) would be permitted to work on (see Section 3).

It is important to account for these factors because **the short-term nature of surge assignments may make divestment** (as a potential way to remedy actual or potential conflicts of interest) **unappealing for most private-sector AI security talent**, who may have significant equity-related interests. Furthermore, **talent may be discouraged by the prospect of post-duty restrictions that apply for significantly longer than their days or weeks of service**, and which could potentially interfere with future work. Nonetheless, conflicts rules undoubtedly serve a valuable purpose of guarding against undue private gain from public service, and are arguably even more critical when individuals maintain ties (via employment and/or equity-related interests) to their outside employer while serving in government—and so might have "divided loyalties." However, the very short duration of surge assignments mitigates the risk, and also leaves less time to remedy conflicts in the midst of an emergency. Moreover, crisis response might be appreciably impaired in its speed and effectiveness without access to these key pools of private-sector AI security talent, making reform that appropriately balances the costs and benefits paramount.

One existing possibility for lessening the conflicts of interest restrictions on IPA assignees (as well as private-sector employees serving in government under any other authority) is treatment as Special Governmental Employees (SGEs), which is available when individuals will only be needed to serve for 130 days or less in any 365-day

[212] Partnership for Public Service, "IPA Agency and Candidate Guidebook," 6.

period.[213] SGE status reduces restrictions to a degree, including by allowing individuals serving in government to continue to receive pay for their service from non-government entities (see Section 3). This is particularly valuable for hiring authorities like expert/consultant and Schedule A, but even for IPA assignees (who may continue receiving their regular salaries from their outside employers pursuant to cost-sharing arrangements with the government), it has benefits. Namely, it allows IPA assignees to retain the proceeds of any accelerated vesting of equity-related interests that they might pursue to remedy conflicts arising after they onboard (on intermittent or "zero-hour" status) and/or when called to active service, once it is clear precisely what kind of work they will need to perform for a particular crisis-response effort.

However, as discussed in Section 3, **SGE status is not a panacea, given that it lessens only some conflicts restrictions.** Moreover, prior recommendations to reduce ethics requirements on IPA assignees targeted nonprofit and academic sending organizations where the lack of profit motives mitigated concerns about vendor lock-in or regulatory capture. In any event, **these ethics concerns are not unique to IPAs**, and in fact affect surge uses of all the hiring authorities outlined in Section 3, given that they all envision private-sector employees serving in government while maintaining their outside employment. Thus, future research should consider how these ethics restrictions might be modified for private-sector AI security talent generally, while balancing real concerns about unfair advantage and other forms of conflicts of interest.

There is recent precedent. Reconciling an expansion of IPA eligibility to for-profit entities and conflicts of interest concerns can be done, as evidenced by the existence and continued operation of the DOD's PPTE.[214] More work should be done to examine how exactly the DOD's PPTE program addresses ethical concerns unique to private-sector employees serving in government, in order to gather valuable lessons that could be applied to a broader expansion of IPA eligibility across federal agencies. To replicate this result, Congress should consider taking legislative action to reform the IPA itself to include the private sector, or alternatively, to create an analogue to DOD's PPTE for crisis-driven surging of AI security talent.

Reforming the Security Clearance Process and Requirements

[213] See also Selinger, "Closing the Science & Technology Gap: Increasing Non-Federal Participation in the Intergovernmental Personnel Act Mobility Program through Ethics Reforms," 474-75 (suggesting that, as the U.S. government "recognizes the concept" of SGEs, it could "reclassify IPAs as something other than federal employees for ethics determinations," page 474). OGE has advised that any IPA detailee serving for 130 days or less in any 365-day period be treated as an SGE for purposes of the criminal conflicts of interest statutes, like the aforementioned prohibition on receiving pay for government service from non-government entities. Office of Government Ethics (OGE), "Memorandum 06 x 10," 3, N.6.
[214] Partnership for Public Service, "Trading Places: The Benefits, Challenges and Potential of Federal Public-Private Talent Exchanges."

Current security clearance processes and requirements may create significant barriers to bringing AI security talent into government. As detailed in Section 3:

- **Clearance investigations are often lengthier than a crisis**—frequently lasting several months—and, if only initiated once an emergency begins, may prevent agencies from accessing outside expertise until long after the crisis has passed. One interviewee noted that DOD recently moved toward accelerating the issuance of *interim* clearances—shortening processing times to days or weeks. However, another expert noted that agencies differ not only in the extent to which they use interim clearances, but also in their practices on how much access these clearances provide, limiting their usefulness.
- **Another major barrier is restrictions on granting security clearances to non-U.S. citizens**, as immigrants make up a large share of the top AI talent in the United States.[215] While immigration reform and expansion of clearance eligibility to non-U.S. citizens are beyond the scope of this paper, future research is recommended to assess the impact of these restrictions on federal AI preparedness.
- Finally, several experts emphasized that **drug-use clearance standards may disqualify or discourage participation from important talent pools**, especially from regions like Silicon Valley, where recreational use is more common. Under current standards, use of illegal drugs—including substances such as marijuana that may be legal at the state but not federal level—within 365 days of a clearance application is often deemed disqualifying.

Two specific reforms might improve the government's ability to grant clearances to the necessary U.S.-citizen talent for service during crisis situations.

First, **agencies should develop expedited security clearance pathways for AI security talent.** While delays in the clearance process affect all candidates in the pipeline, in an AI-related crisis, even a delay of weeks (let alone months) can frustrate the government's ability to contain harm or enact an effective response. Ideally, the government will identify the relevant talent in advance of a crisis and initiate clearances through arrangements such as contracting or hiring for a reserve corps (see above). However, the government should also consider establishing a standing expedited-clearance track for AI security talent that can be activated during emergencies. An expedited clearance pathway ensures that individuals with critical AI security skills are moved to the front of the clearance queue and available for rapid deployment when their expertise is most urgently needed.

As noted in Section 3, the typical cost to an agency for one individual's first-time Top Secret clearance currently starts above $5,000 (and can climb higher when there's a need for priority processing).[216] **Congress and agencies may thus need to allocate more funding for**

---

[215] Oschinski et al., "Strengthening America's AI Workforce."
[216] See Defense Counterintelligence and Security Agency, "Billing Rates & Resources."

clearances, as well as the additional investigative staff likely needed to expedite clearances for AI security talent before and during a crisis. In particular, Congress should appropriate funds for this specific purpose. Barring Congressional appropriations, agencies could direct some of their discretionary funding toward the resources necessary to expedite the clearances of AI security talent for crisis response. Such investment is prudent given that **the price of clearing AI security talent for rapid crisis response could pale in comparison to the economic fallout from a prolonged crisis**.

To complement an expedited clearance pathway, agencies could also further streamline processes, such as running clearance adjudication and any other necessary suitability reviews simultaneously, honoring reciprocity when personnel move between agencies, and making greater use of interim clearances for the relevant experts (and consistently ensuring that such interim clearances allow for meaningful access and work to begin). We recognize security clearance delays can have critical mission impacts across the federal government, and not only limited to AI security talent. However, given the balance of this expertise residing largely outside of government and the speed with which an AI-related crisis can progress, it is important to adapt existing processes to strengthen the government's ability to respond to a crisis without relying on government-wide clearance reform.

Second, **agencies should consider—and coordinate with each other on—systematized guidelines for evaluating past drug use,** with the goal of reducing any automatic disqualifications that would currently exclude key AI security professionals from obtaining clearances. While Congressional action or rulemaking by the Attorney General (via DEA, with an evaluation from HHS) would be required to modify what qualifies as a controlled substance,[217] agencies can consider shortening the "look-back window" over which an individual's past drug use is considered, specifically for highly qualified AI security professionals whose expertise is critical to national security objectives. Experts interviewed for this project discussed how that look-back window varies significantly between agencies, and the adjudicative guidelines issued by ODNI do not specify any particular length for that window, emphasizing instead a holistic "whole-person" approach to clearance decisions.[218] This sort of reform acknowledges that many of the nation's most capable AI security experts may have used substances in the past that could disqualify them from government service, despite their essential skills, willingness to serve, high performance, and access to sensitive information in their home organizations.

Alternatively, agency heads could systematize and make greater use of tools like waivers to exclusions based on recent prior drug use for clearance applicants when specific criteria are met. Such waivers could be available when 1) an individual possesses specialized skills, knowledge, or expertise—documented in writing through something akin to a "letter of compelling need" (identified by one expert as a tool sometimes used in the context of security clearances for contractors)—essential to a specific, time-sensitive U.S. government mission,

[217] See Lampe, "Does the President Have the Power to Legalize Marijuana?," 2-3.
[218] See ODNI, "Security Executive Agent Directive 4 (SEAD-4)," 6 and 17-19.

and 2) when a risk-benefit assessment finds that operational advantages clearly outweigh any risks. Of course, they should not be granted in instances of egregious or other extreme use cases.[219] Experts have indicated that such waivers can already be granted by agencies on an exceptional basis; we recommend formalizing and/or expanding this process to include AI security talent specifically.

The recommendations above seek to alleviate the difficulty of granting clearances to AI security talent specifically, but do not address concerns about the ongoing maintenance of these clearances. We assess that making it easier to grant clearances to AI security experts will improve the government's ability to surge this critical talent in times of need, but leave to future researchers the question of whether it is possible to further widen the pool of willing talent by revising requirements for maintaining clearances as well.

Building the Capacity to Hire

Even if policymakers adopt all of the recommendations above, true "surge" hiring in the midst of a crisis will be difficult to operationalize. Identifying, vetting, and onboarding top talent requires considerable skill, domain expertise, and operational capacity. While today's personnel and processes are frequently strained in meeting existing technical hiring needs, they are even less likely to function at the scale and speed envisioned by this paper. This problem would likely become all the more acute if multiple agencies needed to recruit AI security talent simultaneously, as coordination costs would increase while collective hiring resources would be stretched thin.

Recognizing these important logistical challenges, federal agencies and Congress may wish to take steps to develop agencies' hiring capacity. To begin, as part of the talent inventory discussed earlier in this section, **agencies should assess whether their hiring managers and human resources personnel possess the necessary technical fluency to identify, evaluate, and recruit relevant AI security talent**. In many cases, agencies may find that existing staff lack the domain expertise required to differentiate between highly specialized skill sets, particularly in emerging areas such as model evaluation or adversarial defense. They may also simply lack the relationships and on-the-ground context to know where this talent currently resides and successfully recruit them to government service.

Agencies can then use this inventory of AI "hiring-enabling" capacity to identify and fill gaps in recruitment infrastructure. In the short term, this might look like prioritizing the recruitment of talent managers with expertise in these technical fields and upskilling existing hiring managers, potentially through the use of third-party contractors to provide training. It may also involve developing relationships with key talent pools and developing mechanisms for

[219] See ODNI, "Security Executive Agent Directive 4 (SEAD-4)," Appendix B (describing statutory restrictions, called the Bond Amendments, prohibiting agency heads from granting clearances to "addicts" and active "unlawful users," and barring waivers for such individuals).

interagency coordination, to facilitate hiring that cuts across agency priorities. **Congress can facilitate all of the above by appropriating funds specifically for investment in this AI hiring-enabling capacity**.

Acknowledgements

We thank the following individuals for their input and feedback on our research: Michael Aird, Bill Anderson-Samways, Haydn Belfield, Michael Boyce, Asher Brass, Miles Brundage, Drew Calcagno, Trey Carter, Christian Chung, Shaun Ee, Safa Edwards, Harrison Gietz, John Halstead, Richard Johnson, Laura King, Mike Knapp, Eli Lifland, Sam Marullo, Katrina Mulligan, Dennis Murphy, Kyle O'Brien, Guy Brynt Parmeter, Matteo Pistillo, Andrew Reddie, Faith Savaiano, Aaron Scher, Yonadav Shavit, Matt Smith, Sushmita Sridhar, Raven Stein, Gabrielle Tarini, Madison Van Doren, Hjalmar Wijk, Alexa Wehsener, Peter Wildeford, and Tom Wingfield.

Any errors, omissions, or views expressed in this paper are solely those of the authors. Inclusion in the acknowledgements does not constitute endorsement of our findings or recommendations.

Appendix A: Typology of Talent

Figure 3: Typology of AI Security Talent Potentially Required for Federal Crisis Response

Category	Core Competencies
Foundational Research & Safety Engineering	• Alignment theory, control, and containment • Robust optimization • Interpretability and mechanistic analysis • Safe training pipelines (e.g., RLHF/RLAIF, constitutional RL) • Scalable oversight/reward modeling • Chain-of-thought monitoring • Unlearning
Security Engineering	• Secure model software development lifecycle • Provenance and signing of data/weights • Data center security and infrastructure hardening • Securely storing weights • Access control
Model Evaluation & Red-Teaming	• Frontier evaluation design • Capability testing • Evaluating models against existing benchmarks • Automated exploit/red-team tooling for models • Identification of misaligned model behavior (e.g., deception, power-seeking, unauthorized tool use) • Attribution of harmful outputs to model weights, prompts, or fine-tuning regimes • Immediate incident-level AI forensics (e.g., tracing behavior across versions or instruction sets)
Multi-Agent Interactions & Safety	• Diagnosing and interrupting collusion, feedback spirals, and agent-on-agent attacks • Evaluating emergent collective behavior and circumvention of safeguards • Diagnosing failures of cooperation • Designing defenses such as secure communications protocols and improved network architectures
Engineering & Hardware Implementation & Monitoring	• Inference-time interventions • Real-time monitoring and rollback tools for deployed models • Inference activity tracing and compute usage diagnostics for deployed models • Hardware-level control activation (e.g., location

	verification) • Deployment of chip-level kill-switch design and safe shutdown protocols for model inference
Adversarial ML Defense	• Back-door/poisoning prevention and removal • Safe fine-tuning • Adversarial example research[220] • Model shutdown option design
Application-Layer AI Security & Abuse Prevention	• Prompt-injection and jailbreak resistance • Abuse monitoring and rate-limiting heuristics • Securing model APIs • Safety-as-a-Service tooling
Data Integrity & Forensics	• Poisoning and contamination detection and forensic analysis • Sensitive information detection in large training data corpora • Dataset origin tracing • Data-label integrity and conflict resolution
Systems Safety & Risk Governance	• Evaluating model security controls against existing federal and international AI standards (e.g., NIST's AI Risk Management Framework, ISO 42001) • Formal safety-case authorship

[220]Goodfellow et al., "Attacking machine learning with adversarial examples."

Appendix B: Other Authorities

A few additional federal hiring authorities bear mentioning—and could be subjects of future research or reform with respect to increasing the government's access to AI security talent—although certain attributes make them less relevant at this time to the project of surge hiring for short-term appointments across the government. They are:

1. Schedule B,
2. Direct-Hire Authority (DHA), and
3. Hiring authorities contained in the Defense Production Act (DPA).

First, Schedule B, like the Schedule A authority discussed in Section 3, can be used to hire for positions that are not of a "confidential or policy-determining character" for which it is impracticable to examine candidates.[221] While both Schedules A and B are excepted-service authorities not subject to the full array of ordinary competitive-service hiring procedures, Schedule B is distinct from Schedule A in that, under Schedule B, an applicant must satisfy basic qualification standards established by OPM for the relevant occupation and grade level.[222] The internal agency review necessitated by that requirement may slow down Schedule B hiring relative to Schedule A.

As with Schedule A, OPM must specifically authorize use of Schedule B to a particular end, either on a government-wide or agency-specific basis.[223] In contrast to Schedule A, though, there are few government-wide Schedule B authorizations, none of which appear relevant to AI security talent.[224] That said, there are a number of agency-specific authorizations, such as one granting the State Department authority to hire for scientific and technical positions from among individuals having special qualifications in foreign policy matters, for appointments up to four years in length.[225] OPM could grant the same or similar authority to other agencies (perhaps substituting the State-specific foreign-affairs criteria with other types of agency-specific expertise), giving them the ability to hire across the full gamut of AI security talent.

Second, OPM can grant a Direct-Hire Authority (DHA) to agencies, on either a government-wide or agency-specific basis, to fill vacancies when there is a critical hiring need or severe shortage of candidates.[226] DHA is actually a competitive-service hiring mechanism, versus an excepted-service authority like Schedules A and B, though it shares some of the

---

[221] 5 C.F.R. 213.3201.
[222] Shimabukuro and Staman, "Categories of Federal Civil Service Employment: A Snapshot." Per this report, OPM has in turn instructed agencies to establish these qualification standards.
[223] See 5 C.F.R. 213.3201.
[224] See 5 C.F.R. 213.3202.
[225] See Federal Register, "Excepted Service: Consolidated Listing of Schedules A, B, and C Exceptions," 89 Fed. Reg. 61177.
[226] OPM, "Direct Hire Authority: Fact Sheet."

more streamlined qualities of excepted-service hiring.[227] Specifically, DHA hiring does not entail many of the ordinary competitive-service hiring procedures, like examination and rating and ranking of candidates, though—unlike excepted-service authorities—public notice of vacancies is required.[228] Agencies using a DHA authority must also follow displaced-employee procedures outlined in federal regulations, which generally provide for priority placement of current or former federal employees subject to workforce reductions.[229]

Taken together, such requirements may make the DHA time-to-hire somewhat slower than the excepted-service authorities addressed here.[230] Furthermore, the fact that DHAs can generally be used to hire individuals into permanent career or career-conditional appointments in the competitive service might make hiring managers less inclined to use it for short-term appointments (even though such usage is generally permissible as well), given limitations on funding and personnel bandwidth when making new hires.[231] That said, there are government-wide DHA authorizations in effect covering a variety of IT, cyber, STEM, and even AI-specific roles that could be deployed to increase AI security talent across federal agencies.[232]

Third and finally, the Defense Production Act (DPA) of 1950 includes provisions pertaining to hiring in support of national defense-related emergency response,[233] and so in theory could be quite relevant in an AI-related crisis.[234] Section 703 authorizes an agency head to appoint civilian personnel without regard to the ordinary competitive-service hiring rules, and offers limited flexibility on higher pay for such personnel.[235] Section 710 includes an authorization for the President to create a "nucleus executive reserve for employment in executive positions in Government during periods of national defense emergency, as determined by the President."[236] This "nucleus executive reserve"—typically called the National Defense Executive Reserve (NDER)—was originally conceived as "a pool of executive volunteers," largely from industry,

---

[227] OPM, "Direct Hire Authority: Fact Sheet"; see also 5 U.S.C. 3304.
[228] OPM, "Direct Hire Authority: Fact Sheet." Per that OPM guidance, agencies must post all vacancies to be filled by DHA on USAJobs.
[229] OPM, "Direct Hire Authority: Governmentwide Authority."
[230] However, one expert who consulted on this project said that, in rare cases, DHA time-to-hire could be several days, though that excludes the time needed for security clearances and other processes that must happen prior to onboarding.
[231] OPM, "Direct Hire Authority: Fact Sheet."
[232] OPM, "Direct Hire Authority: Governmentwide Authority." Agencies must have delegated examining authority from OPM in order to use the government-wide DHA authorities, though that covers much of the executive branch. See OPM, "Direct Hire Authority: Fact Sheet"; see also U.S. Merit Systems Protection Board, "Assessing Federal Job Seekers in a Delegated Examining Environment," viii.
[233] 50 U.S.C. 4501 *et seq.*
[234] As of the time of publication, the DPA—save for a few permanent provisions not relevant here—had not yet been reauthorized following its expiration on September 30, 2025. However, reauthorization seems highly likely, as Congress has reauthorized the DPA at least 53 times since 1950. Following its most recent previous lapse (from October 1990 to August 1991), Congress retroactively authorized the DPA to cover the entire period of the lapse. Keys et al., "Reauthorizing the Defense Production Act," 1-2.
[235] 50 U.S.C. 4553.
[236] 50 U.S.C. 4560(e).

"operating in standby mode" until activated during a national emergency.[237] Though this Section 710 authority has remained good law, the NDER has been defunct since the end of the Cold War.[238] Despite efforts to revive it in recent years, no agency currently has an NDER unit.[239]

Executive Order 13603 links sections 703 and 710 together by providing that "[t]he head of each agency with an NDER unit may exercise the authority under section 703 of the [DPA] to employ civilian personnel when activating all or a part of its NDER unit."[240] Thus, the extent to which an agency head might utilize section 703 beyond the context of activating an NDER unit is somewhat unclear. Indeed, publicly available information on when, if ever, agencies have utilized the authority in section 703 is basically non-existent, which could be a function of use in urgent and/or security-sensitive situations—or could indicate that it has simply not been used. Moreover, while section 703 does allow for higher pay relative to the other federal hiring authorities discussed here, the annualized rate tops out at slightly above $200,000, and so is still far from competitive with most private-sector salaries for AI security talent.[241]

The NDER, by comparison, has received more attention over the years, and in fact was the target of reform in a recent legislative proposal.[242] However, as commentators have noted, the NDER proved to be "a weak instrument" even in its Cold War heyday,[243] and "left a bad impression to those familiar with the execution of the program."[244] As documented in a 1983 report, the NDER—which was never activated—suffered from low membership rates, inconsistent recruiting and training efforts by participating agencies, minimal funding, and a lack of specific enough job assignments for members.[245] While some have said that existing authorities are "likely sufficient" as a legal matter for establishing a modern-day NDER, "a major rethinking" of the program and its aims would be necessary to ensure that history doesn't repeat itself.[246]

---

[237] McGinn and Letts, "Options for Strengthening the Use of Defense Production Act Title VII," 24.
[238] See McGinn and Letts, "Options for Strengthening the Use of Defense Production Act Title VII," 6.
[239] See Baker, "A DPA for the 21st Century," 10; McGinn and Letts, "Options for Strengthening the Use of Defense Production Act Title VII," 6.
[240] E.O. 13603 (77 Fed. Reg. 16651), section 501(d), March 16, 2012.
[241] See Baker, "A DPA for the 21st Century," 19 (citing the highest possible rate, circa 2019, as equating to $207,000 per year). Perhaps as another indication of its scant usage, section 703 still includes language capping salaries at the GS-18 rate, which no longer exists. Thus, section 703 salaries would now be calculated according to the Senior Level, which replaced the grades of GS-16 through GS-18. See Baker, "A DPA for the 21st Century," 19; see also OPM, "Fact Sheet: Interpreting References in Law to Maximum Rates Under the General Schedule or 5 U.S.C. 5376."
[242] See H.R. 3561, 119th Congress (FORCE Act).
[243] Fischer et al., "AI Policy Levers: A Review of the U.S. Government's Tools to Shape AI Research, Development, and Deployment," 40.
[244] McGinn and Letts, "Options for Strengthening the Use of Defense Production Act Title VII," 31.
[245] See McGinn and Letts, "Options for Strengthening the Use of Defense Production Act Title VII," 24 and 31 (citing U.S. General Accounting Office, Procurement, Logistics, and Readiness Division, "PLRD-83-51 National Defense Executive Reserve Program.")
[246] McGinn and Letts, "Options for Strengthening the Use of Defense Production Act Title VII," 31.

Appendix C: Additional Information on Methodology for Assessing Federal Hiring Authorities

Who can you hire? First, we assessed whom an agency can hire using a particular authority, e.g., whether an authority is reserved for individuals with scientific or technical expertise, for individuals filling a critical hiring need, etc. Authorities with that sort of focus are an obvious match for surging AI security talent, and so are perhaps less likely to run up against concerns that they overextend existing grants of authority. We also looked at whether there were any restrictions on where you could hire an individual from, e.g., if they must be hired directly from a position in academia. This metric looms particularly large because some authorities—most notably, IPAs—include such limitations, meaning that they cannot be used to draw on expertise from for-profit companies.

How is hiring conducted? Next, we analyzed an array of factors related to how agencies can operationalize a given hiring authority: which agencies can exercise an authority (i.e., whether it is available to all agencies or only to some); whether OPM or any other actor outside an agency must specifically authorize or approve that agency's use of an authority; whether there are any legal limits on the number of hires an agency can make using a particular authority; and how recruitment works.

How long does it take? As an initial matter, we looked at what the average timeline is for use of an authority, generally referred to as "time-to-hire." However, we also gave attention to whether and to what extent timelines for using these expedited authorities might be further compressed during an emergency, if the government were—contrary to the endorsement of advance planning in this report—to neglect to begin hiring until a crisis actually occurred. Specifically, we surveyed the bottlenecks that might make it difficult or infeasible to shorten time-to-hire even in the midst of a crisis, absent prior planning and coordination. These metrics shed light on how broadly and quickly agencies can deploy different hiring authorities, which bears directly on the "surge" aspect of this project.

Who can pay them, and how much? We also examined how pay works for each of the hiring authorities, with special attention to any cost-sharing or other flexibilities regarding how individuals are paid for their time in government service. Such flexibilities would likely make surge hiring AI security talent more cost-effective, and therefore more feasible, for the federal government. Furthermore, retention of normal, outside-of-government salaries could make stints in government more appealing to some people.

What work can and can't they do? We then reviewed several metrics affecting the scope of positions, specifically: whether there are any limits on the type of work that can be performed; whether a position can have a full-time, part-time, or intermittent (as-needed) schedule; and whether there is a cap on the length of the appointment. These metrics reveal important commonalities, as well as noteworthy distinctions. For example, all three of the principal hiring authorities can be used for intermittent appointments, meaning that individuals could be hired, cleared, and onboarded ahead of any crisis, and then called up or "activated" only as needed.That feature is significant given the need to square traditionally long hiring, clearance, and onboarding timeframes with unpredictable, crisis-driven timelines where only short-term appointments are needed. On the other hand, the expert/consultant authority is distinct in containing a prohibition on being used for work performed by an agency's regular employees, meaning that its use could be restricted where agencies already have some in-house personnel working on particular AI-security competencies.

Who is potentially conflicted out? Finally, we considered whether any hiring authorities provided flexibility when it comes to the generally applicable federal conflicts of interest rules, including post-government-employment restrictions. While no such variation exists, the uniformity of these rules across federal hiring authorities highlights the value of SGE status, which lessens some restrictions and so potentially enhances the government's access to key AI security talent who might otherwise be unwilling to serve. It also underscores how other, non-federal-employment mechanisms, such as use of contractors, might factor in as part of surging AI security talent, given that conflicts of interest restrictions operate differently, and generally much less restrictively, in that context.

Bibliography

Advanced Technology International. “Other Transaction Agreements (OTA).” October 2024. https://www.ati.org/ota/.

Advanced Simulation and Computing. “El Capitan: NNSA’s First Exascale Machine.” December 2, 2024. https://asc.llnl.gov/exascale/el-capitan.

Agency for International Development Acquisition Regulation (AIDAR). “Appendix D to Chapter 7—Direct USAID Contracts With a U.S. Citizen or a U.S. Resident Alien for Personal Services Abroad.” Acquisition.Gov., 2025. https://www.acquisition.gov/aidar/appendix-d-chapter-7%E2%80%94direct-usaid-contracts-u.s.-citizen-or-u.s.-resident-alien-personal-services-abroad.

Baker, Jamie. “A DPA for the 21st Century: Securing America’s AI Natural Security Innovation Base.” Center for Security and Emerging Technology, April 2021. https://cset.georgetown.edu/publication/a-dpa-for-the-21st-century/.

Bengio, Y., S. Mindermann, D. Privitera, et al. “International AI Safety Report: The International Scientific Report on the Safety of Advanced AI.” DSIT, January 2025. https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.

Booz Allen. “Inside Booz Allen’s AI Talent Strategy and Workforce.” 2024 https://www.boozallen.com/menu/media-center/q2-2024/inside-booz-allens-ai-talent-strategy-and-workforce.html#:~:text=Data%20indicates%20that%20the%20technology,of%20emerging%20technology%20like%20AI.

Boyd, Aaron. “DoD ‘Force of the Future’ Includes Civilian Digital Service Team.” Federal Times, 2015. https://www.federaltimes.com/management/2015/09/08/dod-force-of-the-future-includes-civilian-digital-service-team/.

Brundage, Miles. “Why Security Comes First.” Miles’s Substack, January 10, 2025. https://milesbrundage.substack.com/p/why-security-comes-first.

Center for Development of Security Excellence. “Personnel Clearances in the National Industrial Security Program (NISP), Version 5.” November 2024. https://www.cdse.edu/Portals/124/Documents/student-guides/IS142-guide.pdf.

Chief Digital and Artificial Intelligence Office. “DoD Digital Workforce.” Accessed September 25, 2025. https://www.ai.mil/Initiatives/DoD-Digital-Workforce/.

Christensen, Michelle D. “Security Clearance Process: Answers to Frequently Asked Questions.” Congress.Gov, October 2023. https://www.congress.gov/crs-product/R43216.

Chuck Grassley. “Grassley Launches Sweeping Review of Program Allowing Privately Employed Individuals to Serve in Federal Government Roles.” January 2023. https://www.grassley.senate.gov/news/news-releases/grassley-launches-sweeping-review-of-program-allowing-privately-employed-individuals-to-serve-in-federal-government-roles.

Cybersecurity and Infrastructure Security Agency. “Cybersecurity Division.” September 24, 2019. https://www.cisa.gov/about/divisions-offices/cybersecurity-division.

DARPA. “AI Next.” September 15, 2018. https://www.darpa.mil/research/programs/ai-next.

———. “SABER: Securing Artificial Intelligence for Battlefield Effective Robustness.” March 12, 2025. https://www.darpa.mil/research/programs/saber-securing-artificial-intelligence.

Defense Acquisition University. “Contracting Cone.” Accessed September 29, 2025. https://aaf.dau.edu/aaf/contracting-cone/.
Defense Civilian Personnel Advisory Service. “Public Private Talent Exchange Program.” 2021 September. https://www.dcpas.osd.mil/learning/broadening/publicprivatetalentexchange.
Defense Counterintelligence and Security Agency. “FAQs – Facility Security Officers.” July 2016. https://www.dcsa.mil/Industrial-Security/Entity-Vetting-Facility-Clearances-FOCI/Facility-Clearances/FAQs-Facility-Security-Officers/.
———. “Billing Rates & Resources.” 2020 https://www.dcsa.mil/Personnel-Vetting/Billing-Rates-Resources/.
———. “Security Assurances for Personnel & Facilities: Limited Access Authorizations (LAA) for Non-U.S. Citizens.” 2022. https://www.dcsa.mil/Industrial-Security/International-Programs/Security-Assurances-for-Personnel-Facilities/.
Delaney, Oscar, Oliver Guest, and Zoe Williams. “Mapping Technical Safety Research at AI Companies.” Institute for AI Policy and Strategy, September 12, 2024. https://www.iaps.ai/research/mapping-technical-safety-research-at-ai-companies.
Donald J. Trump. “Ensuring Accountability and Prioritizing Public Safety in Federal Hiring.” The White House, July 7, 2025. https://www.whitehouse.gov/presidential-actions/2025/07/ensuring-accountability-and-prioritizing-public-safety-in-federal-hiring/.
Donald J. Trump. “Restoring the United States Department of War.” The White House, September 5, 2025. https://www.whitehouse.gov/presidential-actions/2025/09/restoring-the-united-states-department-of-war/.
Doubleday, Justin. “IARPA Looks to next Round of AI Cybersecurity Research.” Federal News Network, April 23, 2025. https://federalnewsnetwork.com/artificial-intelligence/2025/04/iarpa-looks-to-next-round-of-ai-cybersecurity-research/.
Drexel, Bill and Caleb Withers. “Catalyzing Crisis: A Primer on Artificial Intelligence, Catastrophes, and National Security.” CNAS, June, 2024. https://s3.us-east-1.amazonaws.com/files.cnas.org/documents/Catastrophic-AI_TECH-2024_Final.pdf.
Federal Register: The Daily Journal of the United States Government. “Excepted Service; Consolidated Listing of Schedules A, B, and C Exceptions.” 2024. https://www.federalregister.gov/documents/2024/07/30/2024-16724/excepted-service-consolidated-listing-of-schedules-a-b-and-c-exceptions.
Federal Acquisition Regulation. “Subpart 7.5 - Inherently Governmental Functions.” Acquisition.Gov, August 2025. https://www.acquisition.gov/far/subpart-7.5.
———. “37.104 Personal Services Contracts.” Acquisition.Gov, August 2025. https://www.acquisition.gov/far/37.104.
———. “Subpart 9.5 - Organizational and Consultant Conflicts of Interest.” August 2025. https://www.acquisition.gov/far/subpart-9.5.
FEMA. “Advance Contracts of Goods and Services.” September 2025. https://www.fema.gov/businesses-organizations/doing-business/advanced-contracts.
FERC. “Final Report on February 2021 Freeze Underscores Winterization Recommendations.” November 16, 2021.

https://www.ferc.gov/news-events/news/final-report-february-2021-freeze-underscores-winterization-recommendations.
Fiorentino, Dominick A. "Federal Procurement: Emergency-Related Acquisition Flexibilities." March 2023. https://www.congress.gov/crs-product/R47350.
Fiorentino, Dominick A., and Alexandra G. Neenan. "Indefinite Delivery, Indefinite Quantity Contracts." November 2024. https://www.congress.gov/crs-product/IF12558.
Fischer, Sophie-Charlotte, Jade Leung, Markus Anderljung, et al. "AI Policy Levers: A Review of the U.S. Government's Tools to Shape AI Research, Development, and Deployment." University of Oxford, March 2021. https://cdn.governance.ai/AI-Policy-Levers.pdf.
Founders Network. "A Founder's Ultimate Guide to Startup Equity." March 2023. https://foundersnetwork.com/startup-equity/#:~:text=On%20average%2C%20founders%20offer%2010,to%20develop%20an%20ownership%20mentality.
Frazier, Kevin. "Is It Time for an AI Expert Protection Program?" Lawfare, July 2025. https://www.lawfaremedia.org/article/is-it-time-for-an-ai-expert-protection-program.
Gehlhaus, Diana, Ron Hodge, Luke Koslosky, Kayla Goode, and Jonathan Rotner. "The DOD's Hidden Artificial Intelligence Workforce." Center for Security and Emerging Technology, September 2021. https://cset.georgetown.edu/publication/the-dods-hidden-artificial-intelligence-workforce/.
Gilmer, Ellen M. "Homeland Security Department Scraps Cyber, AI Advisory Boards." Bloomberg Government, March 12, 2025. https://news.bgov.com/bloomberg-government-news/homeland-security-department-scraps-cyber-ai-advisory-boards.
Goodfellow, Ian, Nicolas Papernot, Sandy Huang, Yan Duan, Pieter Abbeel, and Jack Clark. "Attacking Machine Learning with Adversarial Examples." OpenAI, February 2017. https://openai.com/index/attacking-machine-learning-with-adversarial-examples/.
Halchin, L. Elaine. "Other Transaction (OT) Authority." Congressional Research Service, July 2011. https://acquisitioninnovation.darpa.mil/docs/GAO%20and%20Congressional%20Reports/CRS%20Report%20OTA%202011.07.15.pdf.
———. "The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues." August 2015. https://www.congress.gov/crs-product/R44157.
Heilweil, Rebecca. "AI Corps Leader Leaves Homeland Security for US Digital Response." Fedscoop, June 16, 2025. https://fedscoop.com/ai-corps-leader-leaves-homeland-security-for-us-digital-response/.
Hirsh, Michael. "The AI Doomsday Machine Is Closer to Reality Than You Think." Politico, September 2, 2025. https://www.politico.com/news/magazine/2025/09/02/pentagon-ai-nuclear-war-00496884.
Homeland Security. "FACT SHEET: DHS Completes First Phase of AI Technology Pilots, Hires New AI Corps Members, Furthers Efforts for Safe and Secure AI Use and Development." October 30, 2024. https://www.dhs.gov/archive/news/2024/10/30/fact-sheet-dhs-completes-first-phase-ai-technology-pilots-hires-new-ai-corps.
———. "Artificial Intelligence Safety and Security Board." May 8, 2025. https://www.dhs.gov/archive/artificial-intelligence-safety-and-security-board.

Hubinger, Evan, Carson Denison, Jesse Mu, et al. "Sleeper Agents: Training Deceptive LLMs That Persist Through Safety Training." Anthropic, January 14, 2024. https://www.anthropic.com/news/sleeper-agents-training-deceptive-llms-that-persist-through-safety-training.

Intelligence Advanced Research Projects Activity. "TrojAI." July 25, 2019. https://www.iarpa.gov/research-programs/trojai.

Keys, Cameron, Adam Levin, and Alexandra Neenan. "Reauthorizing the Defense Production Act." Congress.Gov, June 25, 2025. https://www.congress.gov/crs-product/IN12484

Komitsky, Natalie. "OTA Offers Big Benefits...If You Know How to Use It." Management Concepts, August, 2024. https://www.managementconcepts.com/resource/ota-offers-big-benefitsif-you-know-how-to-use-it/.

Lampe, Joanna R. "Does the President Have the Power to Legalize Marijuana?" Congress.Gov, November 2021. https://www.congress.gov/crs-product/LSB10655.

Langosco, Lauro, Jack Koch, Lee Sharkey, Jacob Pfau, Laurent Orseau, and David Krueger. "Goal Misgeneralization in Deep Reinforcement Learning." arXiv, May 28, 2021. https://arxiv.org/abs/2105.14111.

Los Alamos National Laboratory. "Laboratory Partners with OpenAI to Advance National Security." June 11, 2025. https://www.lanl.gov/media/newsletters/connections/0625-openai-partnership.

Lynch, Aengus, Caleb Larson, and Sören Mindermann. "Agentic Misalignment: How LLMs Could Be Insider Threats." June 20, 2025. https://www.anthropic.com/research/agentic-misalignment.

McGinn, John G. and Olivia Letts. "Options for Strengthening the Use of Defense Production Act Title VII." Acquisition Innovation Research Center, 2025. https://acqirc.org/wp-content/uploads/2025/05/WRT-1097.13.4-v1.4.pdf.

METR. "The Rogue Replication Threat Model." November 12, 2024. https://metr.org/blog/2024-11-12-rogue-replication-threat-model/.

MIT Office of the General Counsel. "MIT Employees Considering Service as a Special Government Employee: Important Topics for Review Beforehand." 2018. https://ogc.mit.edu/latest/mit-employees-considering-service-special-government-employee-important-topics-review.

MITRE. "Symposium on the Federal Workforce for the 21st Century." October 2018. https://www.mitre.org/sites/default/files/publications/PR-18-3746-Symposium-Federal-Workforce-21st-Century-Report.pdf.

Moulange, Richard, Max Langenkamp, Tessa Alexanian, Samuel Curtis, and Morgan Livingston. "Towards Responsible Governance of Biological Design Tools." arXiv, November 2023. https://arxiv.org/abs/2311.15936.

National Initiative for Cybersecurity Careers and Studies. "NICE Framework Competency Areas." August 2025. https://niccs.cisa.gov/tools/nice-framework/competency-area.

National Institute of Standards and Technology. "About BBD." March 13, 2018. https://www.nist.gov/mml/bbd/about-bbd.

———. "Artificial Intelligence Safety Institute Consortium (AISIC)." February 8, 2024. https://www.nist.gov/artificial-intelligence/artificial-intelligence-safety-institute-consortium-aisic.

Neenan, Alexandra G. "Defense Primer: Department of Defense Contractors." June 2024. https://www.congress.gov/crs-product/IF10600.

Nevo, Sella, Dan Lahav, Ajay Karpur, Yogev Bar-On, Henry Alexander Bradley, and Jeff Alstott. "Securing AI Model Weights." Rand, May 30, 2024. https://www.rand.org/pubs/research_reports/RRA2849-1.html.

Nunn, Zachary. "H.R.3561 - FORCE Act." Congress.Gov, May 2025. https://www.congress.gov/bill/119th-congress/house-bill/3561/text.

Office of Government Ethics. "Memorandum 06 x 10." 2006. https://www.oge.gov/Web/OGE.nsf/0/0CC488A084012146852585BA005BED05/$FILE/06x10_.pdf.

Office of Inspector General, Federal Trade Commission. "Audit of the Federal Trade Commission's Unpaid Consultant and Expert Program." August 1, 2022. https://oig.ftc.gov/sites/default/files/reports/2023-08/2022-08-01OIGauditreportunpaidconsultantsFINAL.pdf.

Office of the Director of National Intelligence. "Security Executive Agent Directive 4 (SEAD-4)." 2017. https://www.dni.gov/files/NCSC/documents/Regulations/SEAD-4-Adjudicative-Guidelines-U.pdf.

OpenAI. "Our Structure." May 5, 2025. https://openai.com/our-structure/.

Oschinski, Matthias, Jacob Feldgoise, Jeremy Neufeld, and Sophie Alcorn. "Strengthening America's AI Workforce." March 19, 2025. https://ifp.org/strengthening-americas-ai-workforce/.

Partnership For Public Service. "Roadmap for Renewing Our Federal Government." 2021. https://ourpublicservice.org/our-solutions/roadmap-for-renewing-our-federal-government/workforce/.

———. "Trading Places: The Benefits, Challenges and Potential of Federal Public-Private Talent Exchanges." September 2021. https://ourpublicservice.org/publications/trading-places-the-benefits-challenges-and-potential-of-federal-public-private-talent-exchanges/.

———. "Intergovernmental Personnel Act Hiring and Placement." September 17, 2024. https://ourpublicservice.org/our-solutions/workforce/federal-hiring/intergovernmental-personnel-act-hiring-and-placement/#:~:text=The%20Intergovernmental%20Personnel%20Act%20of,Recruiting%20IPA%20candidates.

———. "IPA Agency and Candidate Guidebook." January 2024. https://ourpublicservice.org/wp-content/uploads/2024/04/IPA-Guidebook-2024-v6.pdf.

Perez, Lisbeth. "DoD Moves CDAO Under Research, Engineering Arm." MeriTalk, August 18, 2025. https://www.meritalk.com/articles/dod-moves-cdao-under-research-engineering-arm/.

Peters, Gary. "Senator Peters Reintroduces Bipartisan Bill to Boost America's AI Workforce Pipeline." April 2025. https://www.peters.senate.gov/newsroom/press-releases/senator-peters-reintroduces-bipartisan-bill-to-boost-americas-ai-workforce-pipeline.

Posard, Marek N., Emily Higgs, Jamie L. Ryan, and Richard S. Girven. "Updating Personnel Vetting and Security Clearance Guidelines for Future Generations." Rand, March 2021. https://www.rand.org/pubs/research_reports/RRA757-1.html.

Princeton University. "Open Faculty & Academic Professional Positions." Accessed September 26, 2025. https://puwebp.princeton.edu/AcadHire/apply/.

Rashidi, Sol. "OpenAI's $10 Million+ AI Consulting Business: Deployment Takes Center Stage." Forbes, July 2025.

https://www.forbes.com/sites/solrashidi/2025/07/16/openais-10m-ai-consulting-business-deployment-takes-center-stage/.
Selinger, Jayme. “Closing the Science & Technology Gap: Increasing Non-Federal Participation in the Intergovernmental Personnel Act Mobility Program through Ethics Reforms.” Public Contract Law Journal 49, July 2020. https://www.americanbar.org/groups/public_contract_law/resources/journal/2020-spring/closing-science-technology-gap-intergovernmental-personnel-act-ethics/.
Shimabukuro, Jon O. and Jennifer A. Staman. “Categories of Federal Civil Service Employment: A Snapshot.” Congress.Gov, March 16, 2019. https://www.congress.gov/crs-product/R45635.
Silva, Felipe Leno da, Ruben Glatt, Brian Giera, et al. “Safety in Artificial Intelligence.” Lawrence Livermore National Laboratory, 2024. https://data-science.llnl.gov/sites/data_science/files/2024-12/SafeAIforDOE%20Digital.pdf.
The White House. “AI Talent Report.” Biden White House, January 14, 2025. https://bidenwhitehouse.archives.gov/cea/written-materials/2025/01/14/ai-talent-report/.
———. “Implementing the President’s “Department of Government Efficiency” Cost Efficiency Initiative.” February 26, 2025. https://www.whitehouse.gov/presidential-actions/2025/02/implementing-the-presidents-department-of-government-efficiency-cost-efficiency-initiative/.
———. “Winning the Race: America’s AI Action Plan.” July 2025. https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf.
United States Government Accountability Office. “Generative AI Use and Management at Federal Agencies.” July 2025. https://www.gao.gov/assets/gao-25-107653.pdf.
U.S. Department of Commerce. “Access to Sensitive Compartmented Information (SCI).” Office of Security. Accessed September 29, 2025. https://www.commerce.gov/osy/programs/personnel-security/access-sensitive-compartmented-information-sci.
U.S. Department of Commerce, National Institute of Standards and Technology. “Glossary: Sensitive Compartmented Information (SCI).” Accessed September 29, 2025. https://csrc.nist.gov/glossary/term/sensitive_compartmented_information.
U.S. Department of Health and Human Services. “HHS Instruction: 304-1: Appointment of Experts and Consultants.” August 17, 2023. https://www.hhs.gov/about/agencies/asa/ohr/hr-library/304-1/index.html#:~:text=As%20of%20the%20date%20of,In%20accordance%20with%205%20U.S.C.
U.S. General Accounting Office, Procurement, Logistics, and Readiness Division. “PLRD-83-51 National Defense Executive Reserve Program.” 1983. https://www.gao.gov/assets/plrd-83-51.pdf.
U.S. General Services Administration. “Alliant 2 Governmentwide Acquisition Contract (GWAC).” 2025. https://www.gsa.gov/technology/it-contract-vehicles-and-purchasing-programs/gwacs/alliant-2.
———. “Alliant 2 Industry Partners.” 2025. https://www.gsa.gov/technology/it-contract-vehicles-and-purchasing-programs/gwacs/alliant-2/alliant-2-industry-partners.

———. “Ways You Can Sell to Government.” 2025. https://www.gsa.gov/sell-to-government/step-1-learn-about-government-contracting/ways-you-can-sell-to-government.
U.S. Government Accountability Office. “Federal Workforce: Opportunities Exist to Improve Data on Selected Groups of Special Government Employees.” July 2016. https://www.gao.gov/assets/d16548.pdf.
———. “Defense Acquisitions: DOD’s Use of Other Transactions for Prototype Projects Has Increased.” November 2019. https://www.gao.gov/assets/gao-20-84.pdf.
———. “Covid-19 Contracting: Observations on Federal Contracting in Response to the Pandemic.” July 2020. https://www.gao.gov/assets/gao-20-632.pdf.
———. “Improved Guidance Could Help Federal Agencies Address Skills Gaps and Maximize Other Benefits.” January 2022. https://www.gao.gov/assets/gao-22-104414.pdf.
———. “Other Transaction Agreements: DOD Can Improve Planning for Consortia Awards.” September 2022. https://www.gao.gov/assets/gao-22-105357.pdf.
———. “Artificial Intelligence: Actions Needed to Improve DOD’s Workforce Management.” December 14, 2023. https://www.gao.gov/products/gao-24-105645.
———. “Innovation Lab.” Accessed September 29, 2025. https://innovations.gao.gov/.
U.S. Merit Systems Protection Board. “Assessing Federal Job-Seekers in a Delegated Examining Environment.” 2001. https://www.mspb.gov/studies/studies/Assessing_Federal_Job_Seekers_in_a_Delegated_Examining_Environment_253638.pdf.
U.S. National Science Foundation. “Artificial Intelligence Scholarship for Service Initiative." May 2024. https://nsf-gov-resources.nsf.gov/files/2024SFSAIReport-r.pdf.
U.S. Office of Government Ethics. “Memorandum Dated February 15, 2000, from Stephen D. Potts, Director, to Designated Agency Ethics Officials, General Counsels and Inspectors General Regarding Summary of Ethical Requirements Applicable to Special Government Employees.” 2000. https://www.oge.gov/Web/oge.nsf/Legal%20Docs/445ECB1FB63809DA852585BA005BED9E/$FILE/00x1.pdf?open.
———. “Ethics Laws Applicable to Special Government Employees.” March 18, 2021. https://extapps2.oge.gov/Training/OGETraining.nsf/%24%24OpenDominoDocument.xsp?documentId=D006291C1FEC02448525869C005BD4B8&action=openDocument.
———. “Conflicts of Interest Considerations: Corporate Employment.” August 2024. https://www.oge.gov/web/oge.nsf/0/EC83872D932E6DCE852585B6005A1F8C/$FILE/Corporate%20Employment.pdf.
U.S. Office of Management and Budget, Office of Federal Procurement Policy. “Publication of the Office of the Office of Federal Procurement Policy (OFPP) Policy.” September 2011. https://prhome.defense.gov/Portals/52/Documents/RFM/TFPRQ/Insource/Docs/OFPP%20Policy%20Letter%2011-01,%20Performance%20of%20Inherently%20Governmental%20and%20Critical%20Functions.pdf.
U.S. Office of Personnel Management. “Fact Sheet: Expert and Consultant Pay.” January 5, 2016. https://www.opm.gov/policy-data-oversight/pay-leave/pay-administration/fact-sheets/expert-and-consultant-pay/.
———. “Direct Hire Authority.” December 29, 2023. https://www.opm.gov/policy-data-oversight/hiring-information/direct-hire-authority/#url=Fact-Sheet.

———. "Direct Hire Authority: Governmentwide Authority." December 2023. https://www.opm.gov/policy-data-oversight/hiring-information/direct-hire-authority/#url=Governmentwide-Authority.
———. "Fact Sheet: Interpreting References in Law to Maximum Rates under the General Schedule or 5 U.S.C. 5376." September 29, 2025. https://www.opm.gov/policy-data-oversight/pay-leave/pay-administration/fact-sheets/interpreting-references-in-law-to-maximum-rates-under-the-general-schedule-or-5-usc-5376/.
———. "Salary Table 2025-GS." 2025. https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/pdf/2025/GS.pdf.
———. "Intergovernment Personnel Act." Accessed September 26, 2025. https://www.opm.gov/policy-data-oversight/hiring-information/intergovernment-personnel-act/.
U.S. Office of Personnel Management, USAJobs. "Volunteers." Accessed September 26, 2025. https://help.usajobs.gov/working-in-government/unique-hiring-paths/students/volunteers#:~:text=Contact%20the%20Personnel/Human%20Resources,for%20specific%20jobs%20or%20functions.
Vadiee, Armani and Todd M. Garland. "The Federal Government's 'Other Transaction' Authority." 2018. https://acquisitioninnovation.darpa.mil/docs/Articles/Briefing%20papers%20THE%20FEDERAL%20GOVERNMENTS%20OTHER%20TRANSACTION%20AUTHORITY.pdf.
Vincent, Brandi. "CDAO's Future Uncertain as Slew of Top Leaders and Tech Staffers Depart." Defense Scoop, May 7, 2025. https://defensescoop.com/2025/05/07/dod-cdao-future-uncertain-top-leaders-tech-staffers-depart/.
Vought, Russell T. "Memorandum for the Heads of Executive Departments and Agencies." Executive Office of the President, April 3, 2025. https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf.
Wasil, Akash R., Corin Katzke, Justin Bullock, and Everett Smith. "AI Emergency Preparedness: Examining the Federal Government's Ability to Detect and Respond to AI-Related National Security Threats." arXiv, July 3, 2024. https://ar5iv.labs.arxiv.org/html/2407.17347.
Wijk, Hjalmar. "Autonomous Replication and Adaptation: An Attempt at a Concrete Danger Threshold." AI Alignment Forum, August 16, 2023. https://www.alignmentforum.org/posts/vERGLBpDE8m5mpT6t/autonomous-replication-and-adaptation-an-attempt-at-a.
Wittmann, Bruce J., Tessa Alexanian, Craig Bartling, et al. "Toward AI-Resilient Screening of Nucleic Acid Synthesis Orders: Process, Results, and Recommendations." bioRxiv, December 4, 2024. https://www.biorxiv.org/content/10.1101/2024.12.02.626439v1.full.
Wuthmann, Walter. "Massachusetts Voters Reject Proposal to Legalize Certain Psychedelic Drugs." NPR, November 2024. https://www.npr.org/2024/11/06/g-s1-33105/massachusetts-legalize-psychedelic-drugs.
Zetter, Kim. *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*. Crown, 2015.