Subcommittee on Cybersecurity and Infrastructure Protection

Committee on Homeland Security

Hearing:

"The AI Security Landscape: How Frontier Models, Agentic AI, and AI Coding Tools Are Reshaping Critical Infrastructure Security and Resilience"

Statement of Matthew Guariglia, Ph.D

Senior Policy Analyst

Electronic Frontier Foundation

June 4, 2026

As a Senior Policy Analyst for the Electronic Frontier Foundation, I thank the Chairs, Ranking Members, and Members of the Committee for the opportunity to share EFF's views on the benefits and harms associated with AI, particularly Generative AI, and options for responsible government interventions in its continued development and deployment. EFF is a nonprofit organization dedicated to protecting privacy, innovation, and free expression in the digital world. EFF is primarily funded by contributions from more than 30,000 dues-paying members. More than 80% of that funding consists of donations under $10,000. We receive less than five percent of our funding from corporate sponsors.

For 35 years, EFF has represented the interests of technology users, both in court and in policy debates, to help ensure that law and technology support and enhance our civil liberties.

AI is an all-encompassing term that seems to grow to include more technologies and use cases by the day. It is vital that we are clear about what we mean when discussing these tools in a cybersecurity context. For example, the use of general-purpose AI models to compile all possible information on citizens from a variety of sources poses a significant threat to privacy. By contrast, the use of a niche model for a specific purpose, such as improving accessibility on a website for the vision-impaired or models tasked specifically with finding vulnerabilities in critical infrastructure, poses less risk to privacy and civil liberties.

AI tools can be useful for all manner of hobbyists, scholars, and businesses, but they can also be misused and misunderstood. And when they become part of critical national security and cybersecurity infrastructure, those risks only increase.

Accordingly, governments must not adopt emerging and powerful technologies without also adopting strong and clear safeguards to protect Constitutional rights. This is of particular urgency because of the demands recently put on Anthropic by the Pentagon to make their technology available for use for all purposes, including those it was not

designed for, like mass surveillance of Americans.[1] EFF opposes the use of generative AI for the purposes of mass government surveillance because that use supercharges unconstitutional violations of civil liberties and because government secrecy prevents the public and lawmakers from knowing when generative AI models make mistakes. That is the baseline upon which all the rest of our recommendations rest.

I. National Security

There is great temptation, motivated by both national interest and lobbying from for-profit enterprises, to deploy emerging technology as quickly as possible. In fact, the Pentagon is already making rapid strides to deploy AI models without the rigorous testing and trial periods done by previous administrations.[2] While the desire to gain a strategic and computational edge is understandable, this rapid deployment raises a number of concerns as to whether the Intelligence Community and the military are meeting their transparency, accountability, and civil liberties obligations.[3] What is worse, this rapid deployment may *also* undermine our national security infrastructure.

A. Supercharging Surveillance

In 2024, the Biden Administration issued a memorandum[4] declaring the intentions of the national security apparatus to leverage the private sector's AI expertise for the public benefit. Unfortunately, meeting that goal may involve merging proprietary and

---

[1] Matthew Guariglia, "*The Department of Defense Wants Less Proof its Software Works*," Electronic Frontier Foundation (October 31, 2025), available at https://www.eff.org/deeplinks/2025/10/department-defense-wants-less-proof-its-software-works

[2] Id.

[3] Matthew Guariglia, "*The U.S. National Security State is Here to Make AI Even Less Transparent and Accountable,*" Electronic Frontier Foundation (November 19, 2024), available at https://www.eff.org/deeplinks/2024/11/us-national-security-state-here-make-ai-even-less-transparent-and-accountable

[4] *"National Security Memorandum on Advancing the United States' Leadership in Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security and Trustworthiness of Artificial Intelligence"* White House (October 24, 2024), available at https://www.presidency.ucsb.edu/documents/national-security-memorandum-advancing-the-united-states-leadership-artificial

government secrecy. This level of opacity creates a system where training data, algorithmic decision-making processes, and use cases related to surveillance and data analytics, including against Americans, are inaccessible to plaintiffs, independent auditors, and anyone whose liberties are impinged by use of AI tools. It also means a system whereby illegal and unethical actions or mistakes and hallucinations enacted by AI might go unknown to the public.

For decades, EFF has fought[5] to expose and challenge secret government interpretations of national security statutes, under which the government has unconstitutionally spied on Americans and retained and analyzed that information. Between open-source intelligence methods like social media surveillance, purchasing intimate data like geolocation from the data broker market, and more traditional methods of signals intelligence which leverages data from telecommunications companies and internet providers, the Intelligence Community is able to access extraordinary amounts of data—not just from surveillance targets overseas but also from Americans who have Constitutional protections from warrantless and indiscriminate surveillance.

Artificial intelligence poses two big, interrelated problems in relation to decades of mass surveillance infrastructure and mostly classified legal interpretations. The first is the drastically increased capacity of AI to analyze information on behalf of the government. Modern AI can synthesize massive amounts of personal data into a comprehensive, exceedingly intimate portrait of an individual's private life, including their political affiliations, religious beliefs, personal communications, medical conditions, and sexual activities. For example, an LLM could infer an individual's association with a particular mosque based upon frequent visits to the mosque's website, engagement with the mosque's social media posts, and their cell phone's physical proximity to the mosque during religious services.

---

[5] Andrew Crocker and Aaron Mackey, "*Victory! EFF Wins National Security Letter Transparency Lawsuit,*" Electronic Frontier Foundation (May 14, 2019), available at https://www.eff.org/deeplinks/2019/05/victory-eff-wins-national-security-letter-transparency-lawsuit?language=el

In a strategic context, one can see many use cases where this would be a beneficial tool. From a civil liberties perspective, however, it becomes more problematic. AI analytic tools can easily combine information that can reveal sensitive personal information about an individual—without requiring preexisting probable cause. Before there was a smart phone in every pocket, many of our security protections relied on the high cost of surveillance. It took quite a lot of individual work to track any one person, and there was no possibility of historical data collection. With an assist from data brokers and government databases, AI has the potential to exponentially increase government capacity to track huge swaths of the population, exposing Americans to granular levels of surveillance with the click of a button.

These risks are already realities in other countries. For example, in China, the government uses AI to analyze and integrate vast amounts of personal information collected through social media monitoring, surveillance cameras, facial recognition systems, and other forms of surveillance. AI can use this information to identify dissent, allowing the Chinese government to locate dissidents and censor government criticism.[6] In a domestic policing context, real-time crime centers and the AI-enhanced platforms that fuse information from sometimes hundreds of sources are also making these fears a reality.[7] Use by the federal government and the data streams that accompany national security signals intelligence, including private digital communications harvested under FISA Section 702, would drastically exaggerate the problem.

The second major concern is the inadequate systems of human oversight. Time constraints, protocols, or disinterest prevent verification of determinations created by AI systems, exacerbate the inherent difficulties in ensuring accountability with a technology that can process large amounts of information. In many contexts, we are already experiencing the ramifications of mistakes in AI analysis and decision-making

---

[6] Darrell M. West, "*How AI Can Enable Public Surveillance*," Brookings Institute (April 15, 2025), available at https://www.brookings.edu/articles/how-ai-can-enable-public-surveillance/.

[7] Andrew Guthrie Ferguson, "*Real-Time Crime Centers and The Brady Puzzle,*" Boston University Law Review (forthcoming 2026), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6468120

processes. Hallucinations, misreadings of inputs or training data, or overstepping guardrails are all common occurrences that are having a range of consequences for people in the United States. While fictional court cases generated by AI[8] may feel like a funny headline, other inaccuracies have had far less benign outcomes. From housing and employment to policing and immigration decisions, AI has already been integrated into many bureaucratic decision-making processes using often unknown or invasive data streams to make impactful choices for a person's life.[9]

"Rubber-stamping" AI decisions or conclusions, combined with inadequate verification procedures, creates a climate where often unaccountable and private AI systems are directing consequential government decisions rather than serving the will and priorities of civil servants. In the national security context, these types of mistakes and the lack of transparency and accountability that follows could have devastating effects on all of our safety.

B. Addressing Cybersecurity Risks

1) Understanding Frontier Models, Agentic AI, and AI Coding Tools

When most people think about AI, they are probably thinking about "generative" AI. Generative AI models—like ChatGPT, Google Gemini, or Claude—are computer systems that can respond to requests, or "prompts," in plain language. The responses can include a plain language response, or images, video, or audio. Or, importantly for the purposes of this Committee, the response can include the source code of a computer program.

---

[8] Pamela Langham, "*Massachusetts Lawyer Sanctioned for AI-Generated Fictitious Case Citations,*" Maryland State Bar Association (March 4, 2024), available at https://www.msba.org/site/site/content/News-and-Publications/News/General-News/Massachusetts_Lawyer-Sanctioned_for_AI_Generated-Fictitious_Cases.aspx

[9] Adam Schwartz and Catalina Sanchez, "*Americans are Uncomfortable with Automated Decision-Making,*" Electronic Frontier Foundation (September 3, 2024), available at https://www.eff.org/deeplinks/2024/08/americans-are-uncomfortable-automated-decision-making

The types of generative AI models that can analyze and produce text or computer code are called large language models, or LLMs. Current LLMs are capable of analyzing complicated questions and producing meaningful responses. The ones that can analyze the most complicated questions and produce the best answers are considered "frontier AI." As the AI industry invests in new models, the "frontier" moves, so an AI model that is considered to be at the frontier today will probably not be at the frontier in a couple of months.

At the time of this hearing, five US companies are generally considered to be operating at the frontier of AI development. There is strong competition from the entire world to develop similar or better AI models. And even those models not on the frontier are remarkably capable.

Some tools go beyond answering questions and are also capable of performing actions related to those questions. This is called "agentic AI," and is a combination of text-generating LLMs with additional software that understands how to perform certain actions. Think of the LLM as a brain in a jar, and an agentic system as a faithful lab assistant that knows how to carry out certain commands on behalf of the brain.

For instance, an LLM by itself can answer the question "read the source code for this program and tell me about all the bugs." An agentic AI system that is configured with an action to send email can both answer questions and take action. For instance, an agentic AI system configured to send emails could take action in response to the prompt "find all the bugs in this source code, and email them to example@example.com."

Agentic AI is particularly effective in the context of AI coding tools. Plain generative AI can write a program, but might make mistakes. An agentic AI coding tool can write a program and also send a command to the operating system to run a compiler on that program. The compiler checks for programming errors. If it produces an error, the agentic AI coding tool can make changes to the program it wrote and try again until the compiler succeeds. If the compiler succeeds, the tool is done and can report success.

Agentic techniques allow AI coding tools to perform more complex tasks with a higher chance of success, without the need for much human interaction. In particular, finding software security vulnerabilities has historically been a high-skill task requiring insight and understanding of the code being studied. Recent advances in both the core models and the agentic "harnesses" they are used with have dramatically reduced the skill required to find vulnerabilities.

2. AI, Cybersecurity, and the Problem of Vulnerability Hoarding

As AI tools have gotten better at writing code, they have also gotten better at finding cybersecurity vulnerabilities in existing code. A vulnerability is a bug that allows an adversary to break a computer system in some way, creating, in turn, an opportunity to crash it, extract private data, or secretly take control of it.

Security experts—both attackers and defenders—search for software vulnerabilities. Attackers do so to exploit them. Defenders do so to fix them so they cannot be exploited in the future. In intelligence and security operations, government acts, at different times, as an attacker and a defender. In espionage, government sometimes exploits vulnerabilities to gain access and extract information. At the same time, government acts to protect its own citizens and critical infrastructure from attackers.

This dual role creates a tension: when a government agency discovers a vulnerability, should the agency hoard it for future exploitation or disclose so it can be fixed? The ubiquity and power of AI vulnerability research decisively tilts the scales in favor of disclosure. For example, the NSA found and developed a number of specific exploits, such as EternalBlue and EpMe/Jian,[10] and kept them open for surveillance purposes, only to have bad actors and adversarial nation states find and use those exploits.[11]

---

[10] "*Eternal Blue*," Wikipedia (last accessed June 1, 2026), available at https://en.wikipedia.org/wiki/EternalBlue

[11] Andy Greenberg, "*China Hijacked an NSA Hacking Tool in 2014 – and Used It for Years,*" Wired (February 22, 2021), available at https://www.wired.com/story/china-nsa-hacking-tool-epme-hijack/

This will happen again. A vulnerability is an observable fact and it is common for the same vulnerability to be independently discovered multiple times. As AI tools make vulnerability research cheaper, independent discovery will become more common and holding vulnerabilities secret will have less value. And at the same time, fixing them promptly will have much more value.

3. How the Government Can Help.

These are areas where government investment could support cybersecurity without intruding on individual rights. First, the government should fully commit to disclosing and helping to fix software bugs, rather than hoarding them for surveillance.

In addition, the government could help address “the patch gap.” When software publishers fix vulnerabilities, users of that software must still update to the latest version in order to receive the benefit. The time between a fix becoming available and a given user installing it is called the patch gap. A brief patch gap—updates installed promptly and reliably—is good for cybersecurity. A long patch gap is disastrous because it gives attackers plenty of time to exploit unpatched systems.

On far too many of our critical and governmental systems, the patch gap is long. Some systems may be entirely unpatchable. But for many systems, the government could support AI-powered security research by making fundamental investments in cybersecurity hygiene: updating all systems promptly, keeping accurate inventories, retiring unmaintained systems and software, and reducing the attack surface of all systems.

The government can also help reduce the national attack surface by investing in the development and deployment of memory safe software in critical systems. Different software development methodologies produce vulnerabilities at different rates. Choice of programming language in particular plays a huge role. Many modern programming languages invest in a property called “memory safety” that automatically detects and prevents certain common bugs. At least 65% of software vulnerabilities are due to the

lack of memory safety[12]. Longstanding programming languages C and C++, notably, do not have memory safety and have no credible path to adding it. In a world where vulnerabilities are easy to find, one of the best defenses is to have fewer vulnerabilities.

II. Other Remedies and Danger Mitigation Options

When the tech companies themselves are trying to insist on guardrails that the military is trying to override, it is up to Congress to step in and provide necessary and balanced regulations, for both the protection and continued security of the American people.[13]

A. Transparency

As noted, the government should share information about newly discovered vulnerabilities, threat models, and mitigation measures with Congress and the public.

Excessive secrecy—often attributable to overclassification and/or overbroad protections for vendors' commercial information—thwarts Congressional oversight and conceals risks to security and civil liberties.[14] Congress should not allow any administration to withhold information it needs to exercise its constitutional prerogatives. Members of Congress should have more access to information about what the Administration has done and is doing.

More broadly, Congress should create a statutory framework for classified information that would restore checks and balances to what has become a system of executive classification. Without greater transparency, Congress will lack the information it needs to identify regulatory gaps and pass carefully targeted laws that address true AI risks

---

[12] Alex Gaynor, *"What Science Can Tell Us About C and C++ Security,"* AlexGaynor.Net (May 27, 2020), available at https://alexgaynor.net/2020/may/27/science-on-memory-unsafety-and-security/

[13] *Supra Note 1*

[14] Faiza Patel and Patrick C. Toomey, *"Bring Transparency to National Security Uses of Artificial Intelligence,"* Just Security (April 4, 2024), available at https://www.justsecurity.org/94113/bringing-transparency-to-national-security-uses-of-artificial-intelligence/

without threatening technological development or the freedom to use tools for beneficial purposes.[15]

Greater transparency would also allow security researchers and other experts to analyze the government's cybersecurity protocols and identify aspects in need of improvement—before our adversaries do. Moreover, U.S. entities, such as defense contractors and other large U.S. companies, may be at risk of similar cyberthreats. Public disclosure of new vulnerabilities and emerging threats would allow the private sector to harden defenses and patch vulnerabilities before attacks occur.[16]

Finally, greater transparency enables better public accountability for potential cybersecurity failures, which can compromise individuals' personal information, jeopardize national security, and otherwise profoundly impact the public interest[17]. Given public distrust of both AI and government surveillance, transparency is essential to preserving trust and legitimacy in the eyes of constituents.[18]

B. Limit AI-Enabled Surveillance of Americans

Minimizing the collection of Americans' personal information and improving procedural safeguards against illegal surveillance are also imperative. That means placing sharp,

---

[15] Tori Noble, Katharine Trendacosta, and Kit Walsh, *"Smart AI Policy Means Examining Its Real Harms and Benefits,"* Electronic Frontier Foundation (February 4, 2026), available at https://www.eff.org/deeplinks/2026/02/smart-ai-policy-means-understanding-its-real-harms-and-benefits

[16] Katharine Megas, Angela Smith et al, *"The Importance of Transparency – Fueling Trust and Security Through Communication,"* Cybersecurity Insights, a NIST Blog (April 3, 2023), available at https://www.nist.gov/blogs/cybersecurity-insights/importance-transparency-fueling-trust-and-security-through ; see also Suzanne Spaulding, *"The Importance of Transparency in Cybersecurity: 2023 Security Summit at the Fortinet Championship,"* Fortinet (September 24, 2023), available at https://www.youtube.com/watch?v=wtCsG1LpRzA

[17] Marisol Cruz Cain et al,*"OMB Action Needed to Address Privacy-Related Gaps in Federal Guidance,"* United States Government Accountability Office (March 2026), available at https://files.gao.gov/reports/GAO-26-107681/index.html?_gl=1*163n46x*_ga*MTc1NDU0OTExNy4xNzgwMTA3Mjc2*_ga_V393SNS3SR*czE3ODAxMDcyNzYkbzEkZzEkdDE3ODAxMDg2MzckajU5JGwwJGgw

[18] Valerie Wirtschafter, *"For AI to make government work better, reduce risk and increase transparency,"* Brookings Institute (January 16, 2025), available at https://www.brookings.edu/articles/for-ai-to-make-government-work-better-reduce-risk-and-increase-transparency/

enforceable limits on government spying, requiring meaningful judicial oversight of surveillance, and passing laws that prohibit the government from bypassing the Fourth Amendment by buying personal information in bulk from data brokers and other private entities.

EFF urges Congress to take two immediate steps: First, enact the Fourth Amendment is Not For Sale Act, which passed the House of Representatives in 2024 by a vote of 219 – 199.[19] This bipartisan bill would prohibit the government from purchasing digital data on individuals it would otherwise need a warrant to collect. This is an especially important body of information that should be minimized in order to prevent it from invasive analysis by AI models in light of the Office of the Director of National Intelligence's creation of a streamlined marketplace for the Intelligence Community to purchase personal information from the data broker industry.

The second important step that Congress could take right now to curb the potential for negative impact of AI use on civil liberties is meaningful reform of mass surveillance authority Section 702 of the Foreign Intelligence Surveillance Act (FISA). Section 702 remains one of the major ways that the National Security Agency compels private companies to hand over the digital communications of an unknown number of U.S. persons—at least several thousand of which have been queried by the Federal Bureau of Investigation without a warrant.

C. Avoid Centralized Government Control Over AI Developers and Models

It is essential to preserve the rights of companies to conduct frontier AI research, freely develop multi-use models for the public, and publicly release new models without the need to ask the government for permission to do so. Extraordinary government interventions, like limits on models that companies may release, are neither appropriate nor necessary to address AI risks. Though AI may pose novel security risks, those risks can be effectively mitigated by hardening insecure government information systems.

---

[19] *Fourth Amendment Is Not For Sale Act, H.R. 4639,"* Congress.Gov (last accessed June 1, 2026), available at https://www.congress.gov/bill/118th-congress/house-bill/4639