PUBLIC SUBMISSION

As of: March 21, 2025
Received: March 14, 2025
Status: [redacted]
Tracking No. m89-extr-hkh2
Comments Due: March 15, 2025
Submission Type: API

Docket: NSF_FRDOC_0001
Recently Posted NSF Rules and Notices.

Comment On: NSF_FRDOC_0001-3479
Request for Information: Development of an Artificial Intelligence Action Plan

Document: NSF_FRDOC_0001-DRAFT-1561
Comment on FR Doc # 2025-02305

Submitter Information

Email: [redacted]
Organization: Wiz, Inc.

General Comment

Please find the attached comments from Wiz in response to this Request for Information on the development of an Artificial Intelligence (AI) Action Plan. We welcome further engagement in support of this effort; please feel free to reach out via the email address provided below with any questions.

Attachments

AI Action Plan Wiz RFI Response

March 14, 2025

Faisal D'Souza
Office of Science and Technology Policy
Executive Office of the President
2415 Eisenhower Avenue
Alexandria, VA 22314

Submitted by email to [redacted]

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

Re: Request for Information (RFI) on the Development of an Artificial Intelligence (AI) Action Plan

Wiz is pleased to provide the below comments as the Trump Administration crafts an Artificial Intelligence (AI) Action Plan.

Our company shares the vision of supporting the widespread adoption and development of AI in a manner that ensures resiliency and creates prosperity for Americans, our nation, and the world. As we discuss below, leveraging AI is a de facto reality today, and it must be acknowledged as the baseline. From there, our government can best determine how to foster the efficiencies and benefits it provides while ensuring our digital infrastructure security is modernized in a manner that supports this new reality and its transformative potential.

We recommend that the AI Action Plan:

1) Recognize the ubiquitous nature of AI technologies and lay out a plan for government adoption of the requisite digital infrastructure needed to foster this novel technology;
2) Promote the implementation of industry-proven practices that ensure the resilience of systems that leverage AI technologies, including the deployment of AI Security Posture Management; and,
3) Accelerate cloud modernization to ensure AI can be utilized in federal service delivery.

About Wiz

Ranking among the fastest growing software companies in history, Wiz understands the transformative nature of technology. Our cloud security platform protects nearly half of the Fortune 100 and many leading cloud-native companies—including those driving the

development and deployment of artificial intelligence. The visibility and threat reduction we provide enables our customers to build AI systems as securely and efficiently as possible, as well as identify 'shadow AI' within environments that can introduce significant security and privacy risks.

Wiz's threat researchers are world-class experts who view the cloud from the perspective of an attacker. Their work is recognized for establishing leading industry practices in AI and cloud. Wiz not only publishes these findings for learning and discussion, but also leverages them to inform our threat platform, which protects over one trillion files across one billion resources.

This threat research, along with our prevalence across the globe, gives us unique insight into both the use and security of AI systems. We hope this perspective is useful as the Administration develops its AI Action Plan.

AI Technologies are Powerful, Ubiquitous, Diverse and Evolving Rapidly

Last month, Wiz released our annual State of AI in the Cloud report. Based on a sample size of hundreds of thousands of public cloud accounts, the report highlights where AI is growing, which new players are emerging, and just how quickly the landscape is shifting.

Most notably, the report shows the speed at which AI adoption occurs. Over 85 percent of organizations are now using some form of AI. For comparison, Kubernetes—a decade-old technology that is widely used to manage modern cloud environments—is only present in 75 percent of cloud environments.

The adoption of self-hosted AI technologies in particular rose significantly, from 42 percent to 75 percent in the past year. Much of this surge is attributed to increased adoption of AI in third-party software, creating additional complexity in understanding risk to the environment.

In addition, businesses are rapidly adopting new technologies as they come online. The release of DeepSeek-R1 prompted a surge in adoption. In early February, Wiz found that roughly seven percent of organizations using self-hosted AI models were using models developed by DeepSeek, which represented a 200 percent increase in the month of January 2025 alone. In the last month, the new DeepSeek model netted around 3.2 million downloads on HuggingFace, a popular AI/ML development, training, and deployment platform. A new version of DeepSeek is reported to be likely to launch as soon as this Spring.

It is notable that rapid adoption of DeepSeek is occurring while privacy and security of the model and its related applications have become a central conversation among technologists and policymakers. DeepSeek's licensing agreement makes it free to use and integrate into existing platforms, whereas many other popular models have more restrictive agreements and cost money to utilize in many situations. Additionally, the model runs efficiently, using fewer resources than many of its competitors. This implies that potential security and risk concerns may not deter adoption when there are perceived efficiency gains. Without visibility and monitoring, significant

supply chain issues could be realized for systems which integrate into such AI systems, as well as the proliferation of 'shadow AI' that goes undetected within sensitive environments.

The ubiquitous nature of AI technology and the rapid adoption of new models tells us, quite simply, that the "horse is out of the barn." It is evident that AI's perceived benefits to efficiency and service delivery are driving private sector adoption, and those returns will only become more substantial as the technology matures.

Therefore, it is imperative that the AI Action Plan is predicated on the fact that AI is a widely used technology that will only become more sophisticated and integrated into our broader digital infrastructure.

Ensure the Resilience of Systems that Leverage AI Technologies

Like any new and rapidly evolving technology, AI requires additional focus from technologists and security professionals as it is integrated into our digital infrastructure. Gartner puts a finer point on mounting AI security threats, asserting that "AI technology usage is increasing risk, and without effective governance and security controls they will have damaging unforeseen impacts on organizations." Wiz agrees. AI's omnipresence, coupled with its relatively young code base, carries serious cybersecurity implications.

Some of the risks from AI being integrated into IT systems may be attributed to the race to bring to market new models, software and features. In February, Wiz discovered a critical vulnerability in DeepSeek. This was a straightforward exposed database in the Chinese-based AI model that allowed broad access to internal data, including chat history, secret keys, backend details, and other highly sensitive information. The exposure could have allowed full database control and potential privilege escalation within the DeepSeek environment, without any authentication or defense mechanism to the outside world.

Wiz responsibly disclosed the vulnerability and the DeepSeek engineering team quickly remediated the exposed database. While a particularly egregious security oversight, it is one in a long line of vulnerabilities, misconfigurations and weaknesses in AI systems and technologies. For instance, in September 2023 Wiz discovered that AI researchers accidentally exposed 38 terabytes of private data — including a disk backup of two employees' workstations. The backup included secrets, private keys, passwords, and over 30,000 internal Microsoft Teams messages.

Likewise, in June 2024 Wiz found a major vulnerability in Ollama, one of the most popular open-source code solutions to integrate AI models into other services through application programming interfaces (APIs). The vulnerability allowed attackers to take over self-hosted AI inference servers, steal or modify AI models, and compromise AI applications. While this API is not intended to be internet-facing code, it is often used that way.

None of these cybersecurity risks stem from the AI model itself or the data weights,[1] two areas which have received significant attention thus far. Rather, these issues can be understood in the context of organizations rapidly adopting and deploying a variety of new AI tools and infrastructure to hone their competitive edge. These tools are often at an early stage of development and lack standardized security features—even basic security functionality like authentication. Additionally, due to their young code base, it is generally easier to find critical software vulnerabilities, making them perfect targets for potential threat actors.

While many of the mistakes Wiz finds are typical cybersecurity concerns, there are also unique aspects that organizations need to pay particular attention to with AI. These characteristics include the 'black box' nature of AI systems making them hard to verify, that models are often a complex combination of application and data, and the lack of consensus standards and consistency in how AI is deployed due to its rapid innovation and the lack of learnings established over time.

Notably, isolation[2] has proven to be an issue with AI technologies. Risk to isolation is particularly problematic, as isolation has served as a key tenet of zero trust architecture, which has been foundational to cybersecurity for the past decade.

For instance, Wiz's Threat Research Team conducted extensive tenant isolation research on some of the most popular AI service providers, including Hugging Face, Replicate and SAP's CoreAI Product. Wiz found these services are more susceptible to tenant isolation vulnerabilities since they allow users to run AI models and applications – which is essentially equivalent to executing arbitrary code that may be malicious.[3] In these specific cases, each organization worked to strengthen the security of their environment; however, as AI infrastructure becomes a staple of business environments, the implications of such weaknesses will continue to become more and more significant.[4]

In September, Wiz published findings regarding a Critical (9.0) NVIDIA AI Common Vulnerabilities and Exposures (CVE). We estimated at the time that over 35% of cloud environments were impacted. The vulnerability enables attackers who control a container image executed by the vulnerable toolkit to escape from that container and gain full access to the underlying host system, posing a serious risk to sensitive data and infrastructure. We coordinated the release with NVIDIA, which issued a patch.

---

[1] A model weight is defined by NIST as a numerical parameter within an AI model that helps determine the model's outputs in response to inputs, per NIST SP 800-218A.

[2] Isolation is the principle of maintaining separation between varying segments of technologies by putting in security policies that govern traffic and access. This ultimately limits risk to the environment.

[3] AI models are effectively applications. Often, these models are packaged in formats that allow arbitrary code execution. Attackers could leverage these circumstances to create malicious models that execute cross-tenant attacks when introducing them in shared environments.

[4] Wiz recommends organizations manage these isolation issues through a cloud isolation review until mitigations prove effective: https://www.wiz.io/blog/genai-tenant-isolation

Our research leads to several recommendations for the AI Action Plan:

- While the federal government has shown significant interest in the security of AI models and weights, it has placed less emphasis on how AI models and technologies interact with the systems and technologies to which they are connected or within which they are integrated. These vulnerabilities stem not from an isolated AI model or dataset, but the context in which they exist within larger systems. We were pleased to see the National Institute of Standards and Technology (NIST) National Cybersecurity Center of Excellence (NCCoE) has prepared a draft concept paper to inform development of a possible Cybersecurity Framework (CSF) Community Profile based on the intersection of cybersecurity and AI. **AI security measures should focus on AI in the context of the systems in which they exist rather than individual elements in isolation.**

- As AI becomes deeply integrated into businesses worldwide—and government systems—**the AI Action Plan should acknowledge these security concerns and direct the development of a government-wide strategy for the adoption of 'best of breed' security practices on-par with those expected of public cloud providers and major infrastructure providers.** This includes technically-verified comprehensive visibility into use of AI technologies (including AI models, training data, and related APIs and services), continuous security posture management of those technologies, and security measures across the AI development lifecycle.

- As organizations rush to adopt AI tools and services from a growing number of startups and providers, it's essential to remember that by doing so, we're entrusting these companies with sensitive data and key roles in our digital operations. Protecting sensitive data and critical service delivery must remain the top priority. It's crucial that security teams work closely with AI engineers to ensure visibility into the architecture, tooling, and models being used, so we can safeguard data and prevent exposure. **AI-focused companies, federal contractors and critical infrastructure providers should be expected to deploy foundational AI cybersecurity measures within both their development and production environments to reduce attack surface.**

Wiz also recommends the AI Action Plan promote incentives or mandate actions to ensure secure AI development and implementation for federal government systems and key critical infrastructure. Recommended actions include:

- Requiring a continuous, comprehensive technically verified inventory of AI services and technologies within AI-enabled systems;

- Ensuring the capability to produce an AI Bill of Materials (AI-BOM)[5] on demand for systems leveraging or integrated with AI;

- Implementing AI Security Posture Management (AI-SPM) tools that allow for continuous monitoring and detection of AI service misconfigurations, as well as enforcement of secure configurations as systems are updated; and,

- Build in secure development processes, techniques and toolsets that enforce AI security best practices across the entire AI development lifecycle. This includes the capability to trace AI risks in environments back to source code repositories and updates.

Federal Cloud Adoption is Key for Federal AI Integration

In January, the Center for Strategic and International Studies Commission on Federal Cloud Policy published a report, "Faster into the Cloud: Accelerating Federal Use of Cloud Services for Security and Efficiency," which called for the acceleration of the adoption of cloud technologies by the federal government.

The report notes that "Cloud computing is essential to AI and other advanced software tools, as it provides the massive processing power and data storage required for training complex models and handling large datasets. It allows AI applications to grow and adapt efficiently. Slow adoption of AI tools relative to the private sector will hamper federal delivery of citizen services."

Cloud environments provide the tools necessary to both rapidly build new environments and ensure better security. When implemented properly, there is better manageability, increased observability and more fulsome documentation. Moreover, APIs make communication between environments more uniform. Suffice it to say, cloud computing creates a more coherent environment to build and secure.

Wiz recommends that the AI Action Plan specifically call for the modernization of federal systems to modern cloud environments. On-premise legacy networks are likely to hamper the adoption of AI and limit the Government's flexibility to keep pace with AI's most modern and secure instances.

Conclusion

Artificial Intelligence has arrived, and the pace of change it brings is only going to increase. As President Trump and this Administration seek to solidify the United States' position as the

[5] An AI bill of materials (AI-BOM) is a complete inventory of all the assets in an organization's AI ecosystem. It documents datasets, models, software, hardware, and dependencies across the entire lifecycle of AI systems—from initial development to deployment and monitoring.

leader in AI, it is critical to ensure the resilience of such systems against criminal and nation-state adversarial actions.

To achieve the Administration's vision of AI that "promote[s] human flourishing, economic competitiveness, and national security," we must promote security practice for the development and deployment of AI. Instituting secure AI development practices, improving tenant isolation, and continuously managing the risk of all elements within a system leveraging AI will create the resiliency necessary to accelerate deployment and innovation.

We hope our recommendations will help as you seek to build a foundation upon which the nation can achieve both rapid and secure AI adoption.