PROTECT AI

BUILDING A SAFER AI-POWERED WORLD

Hearing on

Considering DHS' and CISA's Role in Securing Artificial Intelligence

The Subcommittee on Cybersecurity and Infrastructure Protection

December 12, 2023, at 10:00 a.m.
Cannon House Office Building
Washington, D.C.

Testimony of Ian Swanson
Chief Executive Officer
Protect AI, Inc.

Good morning members of The Subcommittee on Cybersecurity and Infrastructure Protection. I want to start by thanking the Chairman and Ranking Member for hosting this important hearing and inviting me to provide testimony.

My name is Ian Swanson, and I am the CEO of Protect AI. Protect AI is a cybersecurity company for artificial intelligence (AI), that enables organizations to deploy safe and secure AI applications. Previously in my career, I was a worldwide leader of AI/ML at Amazon Web Services and Vice President of Machine Learning at Oracle. Protect AI was founded on the premise that AI security needed dramatic acceleration. When I first started Protect AI, we had to convince industries that the need for security of AI was necessary. Now, industries and governments are openly talking about this need, and shifting the conversation from education of AI security to building security into AI. Against the backdrop of regulation, more front-page headlines on AI/ML security risks, and proliferation of AI/ML enabled tech to deliver business value, the recognition for securing AI/ML applications has never been greater.

AI is the development of computer systems or machines that can perform tasks that typically require human intelligence. These tasks can include things like understanding natural language, recognizing patterns, making decisions, and solving problems. AI encompasses machine learning (ML), which, according to Executive Order 14110 is "a set of techniques that can be used to train AI algorithms to improve performance on a task based on data." A ML model is an engine that can power an AI application and differentiate AI from other types of software code. For many companies and organizations, AI is the vehicle for digital transformation and ML is the powertrain. As such, a secure ML model serves as the cornerstone for a safe AI application, ensuring reliability and security akin to how robust software frameworks and high-grade hardware fortify an organization's technology ecosystem. This ML model, in essence, is an asset as indispensable as any other technology asset, such as databases, cloud computing resources, employee laptops and workstations, and networks. AI/ML assets have numerous challenges in developing, deploying, and maintaining it securely. These include:

- **Limited Transparency in the Operations of AI/ML Applications:** The complex nature of AI/ML algorithms leads to challenges in transparency, making it difficult to perform audits and investigative forensics of these systems.

- **Security Risks in AI/ML's Open Source Assets:** AI/ML technologies often depend on open-source software, which, while fostering innovation, also raises concerns about the security and reliability of these foundational elements.

- **Distinct Security Needs in AI/ML Development Process:** The process of developing AI/ML systems, from data handling to model implementation, presents unique security challenges that differ markedly from traditional software development.

- **Emerging Threats Unique to AI/ML Systems:** AI/ML systems are susceptible to novel forms of cyber threats, such as algorithm tampering and data manipulation, which are fundamentally different from conventional cybersecurity concerns.

- **Educational Gap in AI/ML Security Expertise:** There is a critical need for enhanced training and expertise in AI/ML security. This gap in specialized knowledge can lead to vulnerabilities in crucial AI/ML infrastructures.

Based on my experience and first-hand knowledge, millions of ML models are currently operational nationwide, not only facilitating daily activities but also embedded in mission-critical systems and integrated within our

physical and digital infrastructure. These models have been instrumental for over a decade in areas such as fraud detection in banking, monitoring energy infrastructure, and enhancing cybersecurity defenses through digital forensic analysis. Recognizing and prioritizing the safeguarding of these assets by addressing their unique security vulnerabilities and threats, is vital for this nation and any organization striving to excel in the rapidly advancing field of AI which impacts all elements of the American economy today, and into the future.

US businesses and the United States Government use a significant number of machine learning (ML) models for critical processes, ranging from defense systems to administrative task acceleration. Given the importance of these systems to a safe, functioning government, we pose a critical question: If this committee were to request a comprehensive inventory of all ML models in use in an enterprise or a USG agency, detailing their stages in the life cycle (including experimentation, training, or deployment), the data they process, and the personnel involved (both full time employees, government personnel, and contractors), would any witness, business, or agency be able to furnish a complete and satisfactory response?

Secure AI and ML requires oversight and understanding of an organization's deployments. However, many deployments of AI and ML are highly dispersed and can be heavily reliant on widely used open-source assets integral to the AI/ML lifecycle. This situation potentially sets the stage for a major security vulnerability, akin to the 'SolarWinds incident', posing a substantial threat to national security and interests. The potential impact of such a breach could be enormous and difficult to quantify.

Our intention is not to alarm but to urge this committee and other federal agencies to acknowledge the pervasive presence of AI in existing US business and government technology environments. It is imperative to not only recognize but also safeguard and responsibly manage AI ecosystems. This includes the need for robust mechanisms to identify, secure, and address critical security vulnerabilities within US businesses and the United States Federal Government's AI infrastructures.

Qualcomm[1], McKinsey & Company[2], and PwC[3] have shared analysis that AI can boost the US GDP by trillions of dollars. We must protect AI commensurate with the value it will deliver. To help accomplish this, AI manufacturers and AI consumers alike should be required to see, know, and manage their AI risk:

- **See.** AI/ML systems are fragmented, complex and dynamic. This creates hidden security risks that escape your current application security governance and control policies. Manufacturers and consumers of AI must put in place systems to provide the visibility they need to see threats deep inside their ML systems and AI Applications quickly and easily.

- **Know.** The rapidly evolving adoption of AI/ML adds an entirely new challenge for businesses to ensure their applications are secure and compliant. Safeguarding against a potential "SolarWinds" moment in ML is business critical. Manufacturers and consumers of AI need to know where threats lie in their ML system so they can pinpoint and remediate risk. They must create ML Bill of Materials, scan, and remediate their AI/ML systems, models, and tools for unique and novel vulnerabilities.

[1] Qualcomm: The generative AI economy: Worth up to $7.9T. Available at https://www.qualcomm.com/news/onq/2023/11/the-generative-ai-economy-is-worth-up-to-7-trillion-dollars

[2] McKinsey and Company: The economic potential of generative AI: The next productivity frontier. Available at https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier

[3] PwC: PwC's Global Artificial Intelligence Study: Exploiting the AI Revolution. Available at https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html

- **Manage.** AI/ML security vulnerabilities are difficult to remediate. When operational, technological, and/or reputation security risks are identified that could harm customers, employees, and partners, the business must quickly respond and mitigate them to reduce incident response times. Manufacturers and consumers of AI/ML should create documented policies to help improve security postures, employ incident response management processes, enforce human-in-the-loop checks, and meet existing and future regulatory requirements.

Yes, I believe that the government can help set policies to better secure artificial intelligence. Policies will need to be realistic in what can be accomplished, enforceable, and not shut down innovation or limit innovation to just large AI manufacturers. Against this backdrop, the DHS and CISA play a crucial role in fortifying the security of AI applications.

In the past year, CISA has published two important documents with regard to Securing Artificial Intelligence: "Secure by Design" and the "CISA Roadmap for Artificial Intelligence". The Secure by Design document provides a clear articulation of the "Secure by Design" approach, which is a classic and well understood methodology for software resilience. I applaud the work by CISA and support the three "Secure by Design" software principles that serve as their guidance to AI/ML software manufacturers **1/** Take ownership of customer security outcomes, **2/** Embrace radical transparency and accountability, and **3/** Build organizational structure and leadership to achieve these goals. CISA advancing the "Secure by Design" methodology should help foster widespread adoption. Manufacturers of AI/ML must take ownership for the security of their products and be held responsible, be transparent on security status and risks of their products, and build in technical systems and business processes to ensure security throughout the ML development lifecycle - otherwise known as MLSecOps. While "Secure by Design" and the "CISA Roadmap for Artificial Intelligence" are a good foundation, it can go deeper in providing clear guidance on how to tactically extend the methodology to AI/ML.

I recommend the following 3 starting actions to this committee and other US government organizations, including CISA, when setting policy for secure AI/ML:

1. **Create an MLBOM standard in partnership with NIST and other USG entities.** The development of a Machine Learning Bill of Materials (MLBOM) standard, in partnership with NIST and other U.S. government bodies, is critical to address the unique complexities of AI/ML systems, which are not adequately covered by traditional Software Bill of Materials (SBOM). An MLBOM would provide a more tailored framework, focusing on the specific data, algorithms, and training processes integral to AI/ML, setting it apart from conventional software transparency measures.

2. **Invest in protecting the AI/ML open source software ecosystem.** Per a 2023 study by Synopsis Corporation[4], nearly 80% of AI/ML, Analytics, and Big Data systems use open source software. To protect this, CISA and DHS can mandate and direct other federal agencies to rigorously enforce and adhere to standardized security protocols and best practices for the use and contribution to open source AI/ML software, ensuring a fortified and resilient national cybersecurity posture. The committee should help expand Senate Bill 3050, which includes a proposition and directive on the requirement for AI/ML bug bounty programs in foundational artificial intelligence models being integrated into Department of Defense missions and operations, and be inclusive of all AI/ML assets.

[4] Synopsis Corporation: 2023 Open Source Security and Risk Analysis Report. Available at https://www.synopsys.com/software-integrity/resources/analyst-reports/open-source-security-risk-analysis.html

3. **Continue to enlist feedback and participation from technology startups.** It took a startup in the form of OpenAI to open the eyes of the world to the power and potential of AI. As such, when Congress and other authorities look to regulate AI, it is important to have a broad set of innovative opinions and solutions, and prevent only large enterprises from dominating the conversation, ensuring diverse and forward-thinking perspectives are included in shaping future AI policy and regulation.

In closing and as previously stated, I agree with and support the three principles in CISA's "Secure by Design." However, as mentioned in that document, *"some secure by design practices may need modification to account for AI-specific considerations."* To that end, we realize AI/ML is different from typical software applications and these principles will need to be continuously refined. I welcome the opportunity to propose ideas and solutions that will help drive government and industry adoption of MLSecOps practices, which can be enhanced by new technical standards and sensible governance requirements. I and my company, Protect AI, stands ready to help maintain the global advantage in technologies, economics, and innovations that will ensure the continued leadership of the United States in AI for decades to come.

Thank you, Mr. Chairman, Ranking Member, and the rest of the committee, for the opportunity to discuss this critical topic of security of artificial intelligence. I look forward to your questions.