Written Testimony of:

Sam Rubin
Vice President – Global Operations, Unit 42
Palo Alto Networks

Before the:

Committee on Energy & Commerce
Communications and Technology Subcommittee
United States House of Representatives

Regarding:

"Leveraging AI to Enhance American Communications"

November 14, 2023
10:00 AM

Chairman Latta, Ranking Member Matsui, Chairwoman Rodgers, Ranking Member Pallone, and distinguished members of the committee:

Thank you for the opportunity to testify on how artificial intelligence (AI) already is – and will continue to – enhance cybersecurity defenses. This committee's commitment to exploring the role of AI across every sector of the economy, including information and communications technology, is commendable. My name is Sam Rubin, and I am Vice President and the Global Head of Operations for Unit 42, the threat intelligence and incident response division of Palo Alto Networks.

For those not familiar with Palo Alto Networks, we were founded in 2005 and have since become the world's largest cybersecurity company – protecting businesses and government agencies across more than 150 countries. We support 95 of the Fortune 100, critical infrastructure operators of all shapes and sizes, the U.S. federal government, universities and other educational institutions, and a wide range of state and local partners.

My testimony outlines the cyber threat landscape and how AI is enabling cyber attackers; addresses the benefits of AI-powered cyber defenses that will help us gain the upper hand against these adversaries; and offers considerations for policymakers as they think about whether and how to construct AI guardrails.

At Palo Alto Networks, we have a unique vantage point into the cyber threat landscape. What we see on a daily basis is concerning. Adversaries are growing increasingly sophisticated, and AI will only further amplify the scale and speed of their attacks. However, this sobering backdrop only heightens the importance of fully harnessing the substantial benefits AI offers for cyber defense. Indeed, the demonstrated impact of AI-powered cyber defense is already significant. By leveraging AI, each day we detect 1.5 million unique attacks that were not present the day before. This process of continuous discovery and analysis allows threat detection to stay ahead of the adversary, blocking 8.6 billion total attacks each and every day.

Results like these underscore why Palo Alto Networks firmly believes the risky outcome for society would be to *not* meaningfully leverage AI for cyber defense purposes. AI makes security data actionable for network defenders, giving them real-time visibility across their digital enterprises and the ability to prevent, detect, and respond to cyber attacks quickly. We encourage all entities to embrace the importance of AI for this critical use case and look forward to working with policymakers to further promote its adoption.

Staying Vigilant Against the Evolving Threat Landscape

Cyber adversaries are already leveraging AI to advance their tradecraft and will continue to do so going forward. For example, we see evidence that adversaries are using AI to enhance what we call social engineering attacks – phishing emails designed to lure users to "click the link." Historically, these messages have been littered with typos, making their fraudulent nature relatively easy to detect, but they are becoming more accurate and therefore more believable. Adversaries are now able to generate flawless, mistake-free text, causing click-through rates to skyrocket.

Additionally, bad actors are innovating with AI to accelerate and scale attacks and find new attack vectors. They can now execute numerous simultaneous attacks on one company across multiple vulnerabilities. Adversarial use of AI allows faster lateral movement within networks and more rapid weaponization of reconnaissance data. Going forward, there is the potential for a significant surge in malware variants as the cost of creating customized malware drops substantially.

More recently, we have seen a group called Muddled Libra target the telecommunications sector. Among other tactics, these threat actors frequently use social engineering or text messages to lure employees into providing credentials to gain access to organizations.

None of this should be a surprise. Adversaries are always evolving, with or without AI, and we can never be complacent. As cyber defenders, our mission is to understand and track adversarial capability while relentlessly innovating and deploying best-in-class security tools to stay ahead.

Employing AI for Cyber Defense

Despite the evolving threat landscape, we remain confident that we are well-equipped to combat the cyber incursions of today and tomorrow. AI is, and will continue to be, a game changer to help cyber defenses ward off the crooks, criminals, and nation states that threaten our digital way of life. AI supercharges cyber defenses and helps defenders anticipate, track, and thwart cyber attacks to a degree never seen before.

Our product suite, which spans network security, cloud security, endpoint security, and Security Operations Center (SOC) automation, leverages AI to stay a step ahead of attackers.

We first introduced ML capabilities as part of our malware protection offering 10 years ago and have continued to augment our capabilities with AI tools. In 2020, we launched the industry's first 5G-native security offering to provide 5G network slice security and real-time visibility, prevention, and correlation of 5G user and device threats. We now deploy over 30 products that leverage AI, with more in development.

The continued investment in and integration of AI into cyber defense capabilities is important because it provides defenders a more capable and nimble toolkit to analyze network telemetry – yielding powerful insights that guide deployment of protective measures.

AI-Powered Cyber Defense in Action – Significant Benefits for Defenders

For too long, our community's most precious cyber resources – people – have been inundated with security alerts that require manual triage, forcing them to play an inefficient game of "whack-a-mole," while vulnerabilities remain exposed and critical alerts are missed. Making matters more difficult, this legacy approach often requires defenders to stitch together security data from across dozens of disparate cybersecurity products at the same time, a difficult task often counterproductive to achieving desired cybersecurity outcomes. Organizations find themselves drowning in their own data, struggling to operationalize it. Industry research shows

that over 90% of SOCs are still dependent on manual processes, a sure-fire way to give adversaries the upper hand.

This inefficient, manual posture results in suboptimal Mean Time to Detect and Mean Time to Respond times for security operations teams. As the terms suggest, these metrics provide quantifiable data points for network defenders about how quickly they discover potential security incidents and then how quickly they can contain them. Historically, organizations have struggled to execute against these metrics. In fact, a recent Unit 42 report that analyzed real-world cloud breach incident response cases from 2022 found that security teams take *nearly six days* on average to resolve an alert. In contrast, for the most recent attacks, the average amount of time it takes adversaries to move from compromise to data exfiltration is now *just hours, down from 40+ days as recently as 2021*.

AI-Driven Security Operations Centers

AI-driven SOCs can flip this paradigm and give defenders the upper hand. This technology acts as a force multiplier for cybersecurity professionals to substantially reduce detection and response times.

Early results from deploying this technology for our own company networks have been significant:

- On average, we ingest 36 billion events daily.
- Using AI-driven data analysis, we automatically triage that number down to just eight that require manual analysis.
- We have reduced our Mean Time to Detect to just 10 seconds.
- We have reduced our Mean Time to Respond to just one minute for high priority alerts.

Early customer benefits have been similarly encouraging, with mean response times going from weeks and days to hours and minutes, a fivefold increase in incident close out rate, and a sixfold increase in the average amount of security data ingested and analyzed each day. These dramatic

improvements are critical to stopping threat actors before they can encrypt systems or steal sensitive information.

AI-Driven Vulnerability Discovery

AI is also proving to be a game changer that enables network defenders to discover vulnerabilities across their systems *before* adversaries exploit them. It is often said the internet looks very small to an attacker but massive to a defender. After all, an enterprise that closes 99 of its "digital doors" but leaves one inadvertently open is likely destined for a breach.

Entities of all sizes, public and private, have historically struggled to understand and manage their internet-facing attack surfaces. In fact, we often find that sophisticated enterprises actually have 50% more internet-facing assets than what they were internally monitoring – a visibility gap that gives adversaries the upper hand. To solve this challenge, we employ an AI-powered tool that examines the public-facing internet to discover assets, vulnerabilities, and misconfigurations through the eyes of the adversary. It leverages supervised ML models for this continuous mapping and to prioritize remediation efforts. This visibility puts network defenders back in the driver's seat and dramatically reduces detection and response times.

Capability like this essentially acts as an "early warning system" of sorts – helping defenders prioritize remediation efforts before it is too late. Aggregated insights from this tool provide a helpful baseline of key security trends observed across digital infrastructure, which we publish annually in an Attack Surface Threat Research Report. Recognizing the importance of AI for this purpose, the recently released Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence calls for a joint effort between DHS and DoD to leverage AI for enhanced vulnerability discovery.

From reimagining security operations centers to automating vulnerability discovery, AI-powered cyber tools can have a profound impact on the ability of cyber defenders to safeguard the digital ecosystem. National security, data security, and critical infrastructure resilience will be significantly enhanced.

Maximizing AI's Potential for Cyber Defense

As policymakers in the U.S. consider guardrails governing the development and use of AI, a risk-based and stakeholder-involved approach will help minimize harms without stifling necessary innovation. Palo Alto Networks firmly believes the risky outcome for society is *not* leveraging AI for defensive cyber purposes. We offer the following considerations for policymakers as we look to further encourage the deployment of AI-powered solutions for cyber defense:

Build Upon Flexible Frameworks

The NIST AI Risk Management Framework (RMF) serves as a thoughtfully crafted baseline for understanding AI risk that can serve as the cornerstone for any organization or policymaker. The RMF allows organizations to assess their needs and capabilities against the idiosyncratic circumstances in which they use, develop, or deploy AI systems – evaluating both the risks and benefits of those systems. This flexible framework enables organizations to focus on the unique and discrete risks for each AI use case, while paying particular attention to underlying principles like secure AI development and deployment.

Focus on Securing AI Systems Themselves

In addition to providing a flexible framework that can usefully be adopted by a wide range of enterprises, the NIST AI RMF also emphasizes the necessity of securing AI systems themselves. This emphasis is well-placed. With widespread adoption of generative AI tools and Large Language Models, preventing data leakage is more important than ever. Our research teams are investing significant time and effort to understand how threat actors may compromise AI systems themselves as the vector to attack organizations.

We continue to invest in capabilities to help organizations identify sensitive data, effectively manage user access, and implement robust security measures to protect against internal and

external threats in the age of AI. Moreover, to safeguard against the growing risk of sensitive data leakage to AI apps and Application Programming Interface (APIs), Palo Alto Networks launched a new set of capabilities to secure ChatGPT and other AI apps.

Differentiate Between Use Cases, Impacts, and Data Types

We believe policymakers should employ a risk-based approach when considering AI guardrails that takes into account differences in the use cases, the data processed in those use cases, and the potential resulting impacts on individuals. There are fundamental differences in risk between AI systems that, for example, leverage consumer data to make or facilitate consequential decisions with human impact, as compared with those that leverage security data to ensure the robustness and resilience of networks. When AI is used to make consequential decisions about credit, housing, employment, health care, or insurance, for example, those decisions have a significant impact on individuals. By contrast, when AI is deployed to enable cybersecurity tools to stitch together security data much more quickly and effectively than before, the result is that society as a whole is better protected and all individuals benefit.

We urge policymakers to carefully consider the varied nature of AI use cases to ensure that any new guardrails do not unintentionally inhibit the continued and expanded use of AI-powered tools for cyber defense.

Ensure Disclosure Requirements Do Not Have Unintended Consequences

Palo Alto Networks recognizes that impact assessments and risk management disclosures for AI models are increasingly being proposed to improve AI transparency. We urge policymakers to take into account the potential national security impact when formulating the details of disclosure requirements. For example, public disclosures that require information detailing how network defenders use and train AI systems to secure networks could unintentionally create a roadmap for cyber adversaries to break through those defenses, in turn jeopardizing the underlying security of network and information systems.

Recognize and Promote Cybersecurity as a Privacy Enabler

Palo Alto Networks believes that AI-powered cybersecurity is a great enabler of data privacy. Companies should be encouraged to protect data by implementing robust network and information security practices that help prevent cyber incidents and data breaches from occurring in the first place. Allowing data to be collected, processed, and transferred for cybersecurity purposes ultimately supports both privacy and national security objectives.

People and Partnerships

With AI and automation central to modern cyber defenses, we must educate and train the cyber workforce with the advanced skills required for meaningful jobs that complement technological innovation. This is essential to staying ahead of all cyber threats. To that end, we have been encouraged to see the impact of several initiatives aimed at broadening access to cybersecurity education, including the *Palo Alto Networks Cybersecurity Academy*, which offers free and accessible curricula and hands-on labs to academic institutions from middle school through college. Hands-on experiences with cyber and AI benefit the entire ecosystem as they help to upskill our own workforce as well as that of our customers.

Palo Alto Networks also offers several accelerated onboarding programs to help diversify our workforce, including the *Unit 42 Academy*. As full-time members of our incident response and cyber risk management teams, early-career professionals with both university and military backgrounds spend 15 months developing skills through specialized, instructor-led courses, on-the-job training, and mentorship. We are proud to report that our 2023 class is 80% female.

Partnership is in our DNA at Palo Alto Networks and our collective defense depends upon deepening collaborations between industry and government. We continue to see productive collaboration take place across a range of cybersecurity-focused convening bodies, including CISA's Joint Cyber Defense Collaborative (JCDC), the National Security Telecommunications Advisory Committee (NSTAC), and the Information Technology Sector Coordinating Council (IT-SCC), where we serve as members. We are also an active participant in the DHS ICT Supply

Chain Risk Management Task Force. We maintain robust threat intelligence sharing partnerships with DHS, the Intelligence Community, and across the international community to share technical threat data and collaborate to support government and industry response to significant cyber incidents.

We take our partnership with lawmakers – and this committee – seriously. Please consider Palo Alto Networks a standing resource as you continue to consider cybersecurity and AI policy issues.

Thank you for the opportunity to testify. I look forward to your questions.