Written Testimony of:

Wendi Whitmore
Chief Security Intelligence Officer
Palo Alto Networks

Before the:

Committee on Financial Services

Regarding:

"From Principles to Policy: Enabling 21st Century AI Innovation in Financial Services"

December 10, 2025
10:00 AM

Chairman Hill, Ranking Member Waters, and distinguished members of the committee:

Thank you for the opportunity to participate in today's hearing on AI innovation in financial services. My name is Wendi Whitmore, and I am the Chief Security Intelligence Officer for Palo Alto Networks.

For those not familiar with Palo Alto Networks, we are an American cybersecurity company founded in 2005 that has become a global cybersecurity leader – protecting over 75,000 enterprises in over 150 countries worldwide, including 97 of the Fortune 100, 8 of the largest 10 banks in the U.S., as well as the U.S. federal government and critical infrastructure operators.

As this hearing will reinforce, the enormous promise of AI is undeniable. Palo Alto Networks firmly believes that central to realizing this promise – for the financial services industry and beyond – is being thoughtful in our understanding of the AI-cybersecurity nexus. This requires appreciating 1) the opportunity for AI to turbocharge cyber defense, and 2) the necessity to fortify our AI ecosystem through a *Secure AI by Design* approach.

This perspective is informed from our unique vantage point in our industry. We first introduced machine learning (ML) capabilities as part of our malware protection more than a decade ago and now offer over 30 AI-driven products, with more in development. Investing over $1.8 billion in R&D last year, we combine ML, deep learning, and AI to enable real-time, automated protection across networks, cloud environments, endpoints, and the AI ecosystem.

With financial institutions sitting at the center of the digital economy, the imperative for the sector should be clear: simultaneously embrace AI for cybersecurity, *and* cybersecurity for AI.

The Evolving Cyber Threat Landscape

The rapid development of powerful emerging technologies – from advanced AI to quantum computing – promises incredible opportunities for the sector, but also introduces entirely new classes of complex security risks we must proactively address. While AI will transform financial decision-making and the customer experience, quantum computing's arrival will necessitate a fundamental re-engineering of the cryptography underpinning all global transactions. This is why we must take these proactive security actions now: to responsibly secure the future of finance and unleash the full scale of innovation that these technologies will bring to the sector.

Over the past year, attacks have become faster, more automated, and harder to detect, with time from compromise to data exfiltration shrinking dramatically – now 100x faster than four years ago, and in one in five cases completed within an hour. The financial services sector bears disproportionate risk given the value of its data and interconnected systems, while its attack surface expands through Application Programming Interfaces (API), cloud adoption, and third-party integrations. At the same time, firms contend with evolving regulatory expectations, talent shortages for specialized roles, and the persistent tendency to elevate cybersecurity only after an incident.

Generative and agentic AI intensify these pressures by accelerating every phase of the attack chain and enabling threats such as deepfake-driven fraud, Know Your Customer (KYC) evasion via face-swapping, and tailored spear-phishing campaigns masquerading as communications from trusted financial institutions. Recent Palo Alto Networks research shows that agentic AI can compress what was once a multi-day ransomware campaign into roughly 25 minutes, including reconnaissance, compromise, and data exfiltration. As AI becomes more widely adopted, it will also exponentially expand the digital attack surface and create new vectors for adversaries and cyber-criminals to target – like training data and model environments.

These realities underscore the urgency for financial institutions to modernize defenses with AI-driven security operations that move at machine speed, and embrace Secure AI by Design to harden the expanding AI attack surface and reap the benefits of AI innovation.

Meeting the Moment: Leveraging AI for Cyber Defense

At Palo Alto Networks, we see firsthand how AI-driven cybersecurity is essential to protecting privacy, strengthening national security, and safeguarding our digital way of life. The risky outcome for society would be to *not* meaningfully leverage AI for cyber defense.

Every day, Palo Alto Networks detects up to 8.95 million new attacks. The process of continuous discovery and analysis allows threat detection to stay ahead of the adversary. This real-time awareness of the threat landscape allows our company to block up to 30.9 billion attacks each day. This would not be possible without AI.

We are committed to disrupting the status quo of the cybersecurity industry to simultaneously: 1) deliver transformative cybersecurity outcomes, 2) drive much-needed cost rationalization for network defenders, and 3) eliminate inefficient, manual processes. This innovative spirit will be critical to combatting not just the threats of today, but also the emerging risks – like encryption-breaking quantum computing – of tomorrow.

The Legacy Imbalance

Security Operations Centers (SOCs) are the nerve center of organizational cyber resilience, yet defenders across the financial services sector have historically drowned in alerts and fragmented data. For too long, the community's most precious resource – its people – have been forced into manual triage across dozens of disparate tools, an inefficient model that leaves vulnerabilities exposed, burns out analysts, and degrades core metrics like Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR) – critical indicators that provide quantifiable data points for network defenders about how quickly they discover and contain potential security incidents. The average enterprise SOC ingests data from 83 security solutions across 29 vendors, making effective detection and response nearly impossible under legacy practices.

The consequences are stark: in 75% of breaches, logging existed that should have flagged anomalous behavior, but critical signals were buried and never actioned. Industry research indicates that over 90% of SOCs still rely on manual processes, leaving adversaries with the advantage. Palo Alto Networks reports that teams still take nearly six days on average to resolve cloud breach alerts; meanwhile, sensitive data can be exfiltrated in a matter of hours.

Financial institutions, in particular, are drowning in their own data and struggling to operationalize it, facing an unrelenting “more” – more tools to manage, more devices to monitor, more data to analyze, more attack surface to defend, and more attacks to repel. This imbalance undermines incident response at the very moment when speed and precision are paramount, threatening both national and economic security.

AI-Driven Security Operations Centers

AI-driven SOCs can flip this paradigm and give defenders the upper hand, acting as a force multiplier for cybersecurity professionals to substantially reduce detection and response times. The results from deploying this technology on our own company networks are significant.
On average, we ingest 90 billion events daily, distilling them to 26,000 alerts, and ultimately to a single incident requiring manual investigation.

Financial services institutions are steadily migrating to AI-driven SOC platforms where they benefit from:

- 4x more data ingested with a single, complete source of AI-ready data.
- Over 3,000 AI models with real-time prevention and detection for a full context of threats.
- Automated operations to accelerate SOC workflows resulting in 75% less analyst workload and a 90% reduction in MTTR.

One customer, a global financial markets utility, had several overlapping products that were not talking to each other, leading to cumbersome and inefficient processes and overworked analysts. After deploying an AI-powered solution, their MTTR collapsed from 24 hours to 14 minutes. After deploying this same tool, a global insurance company:

- Prevented 22,831 threats and processed 113,271 threat indicators in less than five seconds, and
- Saved 3,000 hours on an annual basis.

Another large bank leveraging the AI-powered SOC saw similarly transformative outcomes, tremendous efficiency gains, and reduced SOC staff attrition:

- 180 hours/year saved by automating SIEM – the system that collects and analyzes log data from across an organization’s IT environment – and threat intel reporting;
- 500 hours/year saved by automated data collection and enrichment;
- 360 hours/year saved by automating 4 CTO playbooks;
- 240 hours/year saved with automated threat intelligence and case information enrichment.

These dramatic improvements are critical to stopping threat actors before they can encrypt systems or steal sensitive information – which is now frequently happening in mere hours. Simply put, none of this would be possible without the power of AI.

AI Security Accelerates AI Innovation

AI adoption is integral to America's innovation leadership, which is why Palo Alto Networks was proud to support America's AI Action Plan.

While AI offers significant benefits to financial institutions, as demonstrated above, its rapid growth largely outpaces the adoption of security measures designed to protect it and expands the attack surface. According to an October 2025 survey from the Conference Board, nearly three-quarters of S&P 500 companies now flag AI as a material risk in their public disclosures, which is a jump from just 12% in 2023. And cybersecurity risk, specifically, was among the most commonly cited risks across all public disclosures.

Traditional security tools rely on static rules that are no match for today's dynamic AI risks. They are easily overwhelmed by alerts and often miss advanced attacks like multi-step prompt injections or adversarial manipulations. The rise of autonomous AI agents makes things even more challenging, as these agents can take unpredictable actions that are difficult to monitor with legacy methods. Most current tools are simply not built to inspect or adapt to these behaviors in real time.

Rapid AI adoption has dramatically expanded the attack surface by exposing organizations' AI ecosystems (the infrastructure, data, models, applications, and agents) to unique threats that legacy cybersecurity solutions were not explicitly designed to address. Unlike traditional cyber exploits that target software vulnerabilities, AI-specific attacks can manipulate the very foundation of how an AI system learns and operates. These attacks are not just about breaching a network but also can be about corrupting the AI's probabilistic logic itself.

It is critical to recognize and address this security gap with defenses that are truly AI-native. AI security must be a cross-cutting imperative across and within every layer of the AI technology stack. By securing AI capabilities with equally advanced cybersecurity technologies, financial institutions can confidently accelerate safe AI adoption.

A Call to Action

Recognizing this new risk paradigm, we are encouraged by recent policy developments that reinforce that AI adoption and AI security can, and must, go hand in hand. American's AI Action Plan calls for "Secure-By-Design AI Technologies and Applications." It further notes, *"The U.S. government has a responsibility to ensure the AI systems it relies on – particularly for national security applications – are protected against spurious or malicious inputs. While much work has been done to advance the field of AI Assurance, promoting resilient and secure AI development and deployment should be a core activity of the U.S. government."*

Voluntary standards bodies, like NIST, are starting to weigh in, initiating work to develop control overlays for securing AI systems or AI security-specific profiles to established cybersecurity risk management standards, like the NIST Cybersecurity Framework.

We are also encouraged by recent federal actions to operationalize Secure AI by Design through acquisition. Most recently, the General Services Administration announced a governmentwide OneGov agreement with Palo Alto Networks to help federal agencies secure their AI deployments through a cost-effective and scalable platform.

Secure AI by Design Policy Roadmap

While the consensus on rapid AI adoption is broad, organizations frequently struggle with the lack of clarity on what effective AI security looks like in practice. To help close the gap between high-level intent and actionable strategy, we developed a Secure AI by Design Policy Roadmap. The roadmap lays out a four-part construct for holistically understanding the unique attributes of the AI attack surface and threat environment (the "What"), ultimately guiding the development of purpose-built, actionable policy strategies and solutions for securing the AI era (the "How").

Secure AI by Design is a comprehensive strategy that empowers organizations to unleash their full potential for AI innovation by fundamentally integrating security throughout the entire AI lifecycle and across the AI ecosystem. This proactive stance ensures that security is a feature, not an afterthought, which is crucial for building trust, maintaining compliance, and mitigating risks.

By embedding security controls from the start, Secure AI by Design provides the visibility, control, and protection necessary to confidently build and deploy AI-powered applications at enterprise scale. Key components of this portfolio, such as AI access security, security posture management, and AI runtime protection, work together to safeguard sensitive training data, secure the AI supply chain, and protect running models against novel threats. This unified, end-to-end security solution allows organizations to focus on the business benefits of AI – like efficiency and innovation – without compromising the confidentiality, integrity, and availability of their AI systems.

The Secure AI by Design approach helps address the four security imperatives organizations most pressingly face in AI adoption:

1. **Secure the use of external AI tools:** Organizations must continuously discover and inventory AI tools, applications, and agents used across the enterprise – including "shadow AI" accessed via browsers or APIs – to govern model risk through scanning and policy, and protect usage through continuous monitoring for emerging attack vectors such as prompt injection, model extraction, and data exfiltration.
2. **Secure underlying AI infrastructure and data:** Organizations must continuously assess security, safety, and compliance risks across infrastructure, applications, models, and datasets – both in the supply chain and at runtime – through automated posture

management, model scanning for malicious or anomalous components and backdoors, and continuous adversarial testing.

3. **Safely build and deploy AI applications:** Security must be integrated into development and deployment via red teaming tailored to AI systems, configuration baselines for inference-time controls and safe tool invocation, guardrails for transactions, and mechanisms to block unsafe or risky models and patterns in production.
4. **Monitor and control AI agents:** As autonomous and semi-autonomous agents proliferate, security must center on identity-first controls, least-privilege access, scoped tool permissions, separation of duties, and governance over agent-to-agent and agent-to-system communications, alongside runtime monitoring for anomalous behaviors.

Leveraging Secure AI by Design imperatives will help to build a framework where security controls and governance are key tenets of AI innovation. Securing the entire AI ecosystem from development through deployment and use – not added as an afterthought – will best position the financial services sector to address the complex threats emerging from rapid AI adoption and fully embrace the benefits of AI.

Secure AI by Design must be anchored in enterprise governance and clear lines of accountability. Financial institutions should maintain risk-tiered AI inventories, document model lineage and intended use, identify and enforce strict access controls to sensitive data, and implement testing and monitoring commensurate with risk. Governance structures, centered on Secure AI by Design, should enable board and senior management oversight, ensure auditable logs across the AI lifecycle, and align with established model risk practices familiar to prudentially regulated firms.

Recommendations to Drive AI Innovation in Financial Services

Palo Alto Networks is proud to be an integrated national security partner with the federal government and stands ready to help. To that end, we developed a set of recommendations for policymakers to consider at this pivotal moment for our nation's cyber defense and AI leadership:

1. Promote AI-driven security operations. Encourage modernization of security operations across the financial services sector by recognizing AI-driven SOC platforms as a best practice to improve detection, response, and resilience. Supervisory guidance should emphasize measurable outcomes – such as reductions in MTTD and MTTR.
2. Champion voluntary Secure AI by Design frameworks. Embed Secure AI by Design across the financial services ecosystem to secure the use of external AI tools, secure underlying infrastructure and data, monitor and control AI agents, and safely build and deploy AI apps.
3. Ensure that AI policies safeguard cybersecurity and innovation. Policymakers should carefully tailor any new guardrails to ensure they don't unintentionally restrict the use of AI-powered tools for cyber defense. Regulatory proposals should be risk-based and

explicitly protect the ability of AI developers or deployers to ensure, maintain, and improve network and information security or to prevent, detect, protect against, or respond to cybersecurity threats.

4. Enable secure innovation through controlled experimentation. Establish AI innovation labs and regulatory sandboxes across the prudential and market regulators to allow banks, insurers, payment companies, and fintechs to test AI systems – including cybersecurity for AI – in controlled environments. Applications should include clear risk management plans covering data protection – especially intellectual property and proprietary information, model robustness testing, and misuse prevention, alongside success criteria that quantify security outcomes and operational benefits.
5. Strengthen public-private collaboration on AI security best practices and AI-enabled threats. Encourage regular, operationally focused exchanges between financial institutions, cybersecurity providers, and regulators to share indicators, detection methods, red team findings, and best practices for securing AI models and infrastructure. Focus particular attention on AI-enabled fraud, identity abuse, and deepfake-driven social engineering targeting payments and high-value transactions.

Building the AI-Ready Ecosystem: People and Partnerships

AI innovation and secure deployment requires people, processes, and technology working in concert. As a Financial Services Information Sharing and Analysis Center (FS-ISAC) sector advisor, Palo Alto Networks contributes to a number of timely reports analyzing the current threat landscape and trends facing financial institutions and actively supports the work of the AI Risk Working Group as it explores AI in cybersecurity and the cybersecurity of AI. Palo Alto Networks is also proud to be a founding Alliance member of CISA's Joint Cyber Defense Collaborative (JCDC), and one of the original private sector partners in the JCDC's AI working group, which developed CISA's AI Cybersecurity Collaboration Playbook.

As a member of the Cyber Risk Institute (CRI) Innovator Program, Palo Alto Networks helps member firms implement the Financial Services Cybersecurity Profile. This harmonized compliance framework aligns with major standards, including NIST Cybersecurity Framework 2.0, the NIST Ransomware Framework, CISA's Cross-Sector Cyber Performance Goals, and SEC disclosure rules. By acting as a common global baseline for regulator examinations, this framework allows financial institutions to optimize compliance resources, reduce reconciliation time for exam issues, and simplify security oversight.

With AI and automation central to modern cyber defenses, it is critical we educate and train the cyber workforce of tomorrow with the advanced skills required for meaningful jobs that complement technological innovation. To that end, Palo Alto Networks is a proud signer of the Pledge to America's Youth. Led by the great work of the Palo Alto Networks Cybersecurity Academy, we are integrating AI into our core curriculum and offering hands-on, AI-in-action labs.

Thank you for your leadership on this issue and for the opportunity to testify. I look forward to your questions.