House Homeland Security Committee

Subcommittee on Cybersecurity and Infrastructure Protection
Subcommittee on Oversight, Investigations, and Accountability

Testimony for the hearing
"The Quantum, AI, and Cloud Landscape: Examining Opportunities, Vulnerabilities, and the Future of Cybersecurity”

Written Testimony of Michael Coates
Founding Partner, Seven Hill Ventures

December 17, 2025

Chairman Ogles, Ranking Member Swalwell, Chairman Brecheen, and Ranking Member Thanedar, I thank you for the opportunity to testify before you today. I'm honored to be here to speak about the changing landscape in cybersecurity and the resulting impacts from AI and quantum computing.

The perspective I will share is grounded in over twenty years of experience in cybersecurity, including service as a chief information security officer, a chairman of a global non-profit advancing the state of application and coding security, a technology startup founder, and a venture capital investor supporting cybersecurity innovation.

Today we sit at the precipice of significant change. While advancements in AI and development towards AGI are widely discussed, the practical and operational impacts to cybersecurity defenders are less often examined.

The fundamental reality is not that AI and quantum are creating new types of threats, but rather they are collapsing the time, cost and skill required to conduct cyber operations. These changes are outpacing the existing technical, regulatory and operational defenses. This shift reshapes the cyber threat landscape and forces a reconsideration of how we defend critical systems in an era defined by speed, automation, and intelligent scale.

What Is Changing: The Compression of Cyber Capability

Capability Compression & Orchestration Expands the Attacker Base

Corporations and citizens potentially face a variety of threat agents including highly funded nation state adversaries, financially motivated cybercriminal organizations, and lone hacktivists motivated by ideology. Each attacker type has different skills and resources at their disposal and to date, these have constrained the complexity or scale of cyberattacks available to each adversary.

The most advanced attacks were often only launched by nation state adversaries against select targets. Whereas cybercriminal entities focused their efforts on pipelines of optimized offensive security services, such as ransomware extortion, to monetize the compromise of businesses or individuals.

Robust security attacks require a series of steps spanning reconnaissance, exploitation, command and control, and delivery of the ultimate objective, such as data theft or system modification. Each of these components could be executed by a well-funded nation state adversary or a competent cybercriminal organization, but it was not as

achievable for the lone hacktivist or unsophisticated security hacker. This is rapidly changing.

As demonstrated in the November, 2025 Anthropic report “Disrupting the first reported AI-orchestrated cyber espionage campaign”[1], a nation state adversary used AI systems as a central brain and point of coordination for a complete security attack against multiple targets across the United States. AI was used to execute and interpret results for each step of the attack and as an overall orchestration layer, with the human adversary only interacting at a few decision points.
While this attack may not have demonstrated new or novel attack methods, the orchestration and use of AI is a critical development in the ecosystem of the cybersecurity adversary.

Agentic Attacks Remove Human Constraints

Agentic AI systems will enable the attacker to no longer be bound by time of day, hours awake, or the need for food or sleep. Autonomous agentic systems are replicating the most advanced attackers and will be able to target with accuracy and ease.
This is no longer theoretical as research just released by Stanford[2] shows that an autonomous AI penetration-testing agent already performs at or above the level of most highly skilled professional security testers, outperforming nine out of ten participants in a live network test with an 82% valid vulnerability discovery rate at a fraction of the cost.

Acceleration of Vulnerability Discovery and Exploitation

Furthermore, the increasing power of AI for software vulnerability analysis is enabling faster and more accurate detection of previously unknown zero day security vulnerabilities. For example, Google’s Big Sleep, a collaboration between Google Project Zero and Google DeepMind, has discovered a critical zero day vulnerability in the major software SQLite Database Engine.[3]

Over the past decades, the challenge for many organizations has not been knowledge that a vulnerability existed, but rather the operational inertia to deploy, test, and productize the software patch. In fact, the 2025 Verizon Data Breach Investigations Report found that vulnerability exploitation was the initial access vector in 20% of breaches, and that defenders often cannot remediate fast enough—organizations fully

[1] https://www.anthropic.com/news/disrupting-AI-espionage
[2] https://arxiv.org/pdf/2512.09882
[3] https://cloud.google.com/blog/products/identity-security/cloud-ciso-perspectives-our-big-sleep-agent-makes-big-leap

remediated only about 54% of vulnerabilities in network edge devices, with a median remediation time of 32 days, while CISA KEV vulnerabilities can be mass exploited in a median of five days.[4]

The Practical Result: Reduced Time for Defenders

With AI orchestration, the ease of launching comprehensive cybersecurity attacks against any target is substantially reduced. The result is that many more potential adversaries now have the means to execute these attacks.

In addition to an increase in attacks against the most critical targets, this development will also result in lesser-profile targets, such as small businesses across the country, being subjected to full-scale security assaults.

The direct result of this change will be a dramatic drop in the time available for defenders to detect attacks, initial compromise, or lateral movement before critical access or sensitive data is breached. Taken together, these shifts do not just increase cyber risk, they fundamentally change the speed at which cyber incidents unfold.

Why Time Compression Changes the Nature of Cyber Risk

The compression of time, cost, and skill required to conduct cyber operations fundamentally changes how cyber risk manifests in practice. While individual techniques may appear familiar, the speed at which attacks now unfold alters the balance between attackers and defenders in ways that existing security models were not designed to accommodate.

The most immediate consequence is a dramatic reduction in the time available for defenders to detect and respond to malicious activity. AI-enabled orchestration and automation allow attackers to move from initial access to lateral movement and impact far more quickly than in the past. In many cases, defenders are no longer responding to early indicators of compromise, but to attacks that are already well underway.

This compression of time disproportionately affects organizations that lack large, specialized security teams. While highly resourced enterprises may be able to invest in advanced detection and response capabilities, smaller organizations, including hospitals, schools, food processing facilities, and small businesses often rely on delayed or manual processes. As sophisticated attacks become easier to launch and

[4] https://www.verizon.com/business/resources/reports/dbir/

less expensive to operate, these lower-profile targets increasingly face the same level of adversarial capability once reserved for critical national infrastructure.

At the same time, intelligent automation and scaling by adversaries is shifting the risk of attacks from periodic events to a continuous threat. AI-driven attacks do not require sustained human attention and can operate persistently, adapting to defenses and retrying failed approaches automatically. This erodes traditional assumptions that organizations can recover between incidents or rely on periodic assessments to maintain security.

Existing defensive and governance models further compound this challenge. Over the past decades, many major breaches did not occur because vulnerabilities were unknown, but because organizations were unable to deploy patches or mitigations quickly enough. As AI accelerates vulnerability discovery and exploitation, this operational inertia becomes more consequential. The gap between awareness and action grows more dangerous as attack timelines compress.

The result is a widening gap between the speed and accessibility of modern cyberattacks and the ability of most organizations to respond. As AI compresses attack timelines and expands the pool of capable adversaries, cybersecurity outcomes will increasingly be determined by whether defenses can operate at machine speed.

Implications for Cyber Defense, Policy and Coordination

The advancements in artificial intelligence and quantum computing present significant opportunities for innovation, but without appropriate alignment between technology, operations, and governance, they also introduce material cybersecurity risk. The shifts described earlier are not theoretical, and they cannot be addressed by any single organization or sector acting alone.

The following are key areas where attention is warranted to increase the cybersecurity posture of our organizations and critical systems.

- **Secure by Design as a Baseline Expectation**

  As software is increasingly written, analyzed, and modified by AI systems, secure design principles must be integrated into the creation of software from the outset. Initiatives such as CISA's Secure by Design program, along with industry standards promoted by organizations like OWASP and the Cloud Security Alliance, provide important guidance. Supporting these organizations and

reinforcing these efforts helps ensure that speed and automation do not come at the expense of security fundamentals.

- **Regulatory Clarity That Supports Speed and Innovation**

Clear and transparent regulatory frameworks are necessary to enable rapid innovation while maintaining responsibility for security and safety. In an environment where threats evolve quickly, ambiguity or fragmentation in regulation can unintentionally slow defensive response and increase systemic risk. Policy should seek to provide clarity and consistency without constraining the ability of organizations to adapt at machine speed.

- **Public–Private Coordination on AI-Driven Cyber Threats**

The pace of change in the cyber threat landscape reinforces the importance of strong public–private partnerships. Effective coordination, information sharing, and joint response mechanisms help ensure that defensive learning keeps pace with adversarial innovation. These partnerships remain a critical component of national cyber resilience as AI-driven threats continue to evolve.

- **Migration Toward Autonomous Defensive Capabilities**

As attackers increasingly rely on automation and agentic systems, purely human-driven defenses will struggle to keep pace. Continued investment in research, development, and deployment of intelligent and autonomous defensive systems is necessary to address machine-speed threats. This includes supporting innovation across both the public and private sectors.

- **Quantum Preparedness for Cryptographic Systems**

Stable, cryptographically relevant quantum computing would render many of today's widely deployed public-key encryption algorithms ineffective, impacting secure communications across government, industry, and critical infrastructure. While post-quantum cryptographic standards already exist, the primary challenge is the time and coordination required to migrate existing systems. Deliberate preparation is crucial to avoid a reality where an adversary achieves cryptographically relevant quantum capabilities first and thus access not only to future communications, but potentially to sensitive data captured and stored today.

- **Trustworthiness and Transparency in AI Systems**

  As AI systems are increasingly embedded into security-sensitive workflows, trust in operation becomes crucial. Large language models reflect the data, incentives, and governance structures under which they are trained, and these factors can materially influence reliability and security outcomes.

  Bias in AI systems — whether intentional or unintentional — can affect how software is generated, how alerts are prioritized, and how decisions are made. In security-critical contexts, performance alone is not sufficient; the provenance, training, and oversight of AI systems must also be considered as part of risk assessment.

  Furthermore, greater transparency in software procurement and composition is needed. Requiring bill of materials and software contracts to disclose the use of AI within software, as well as the specific models and model origins, can help organizations better assess risk and make informed security decisions, particularly in sensitive or critical environments.

Artificial intelligence and quantum computing are accelerating dynamics that dramatically shift the cybersecurity landscape. As AI and quantum computing continue to advance and are increasingly leveraged by cyber adversaries, success will depend on whether our technical, operational, and institutional responses can adapt at comparable pace.

I appreciate the opportunity to share these observations and look forward to your questions.