

**House Homeland Security Committee**  
**Subcommittee on Cybersecurity and Infrastructure Protection**

**“Considering DHS’ and CISA’s Role in Securing Artificial Intelligence”**

December 12, 2023

**Written Statement of**

Alex Stamos  
*Chief Trust Officer*  
*SentinelOne<sup>1</sup>*

Chairman Garbarino, Ranking Member Swalwell, and Members of the Subcommittee, thank you for having me here today to discuss the challenges and opportunities presented by artificial intelligence and machine learning. These world-changing technologies have the potential to impact nearly every aspect of our lives, and they are likely to continue to scale at lightning speeds. This subcommittee, and policymakers at all levels of government, face the challenging task of matching the pace of innovation with thoughtful policies to harness the positive aspects of AI while minimizing its dangers. In the context of cybersecurity, AI and machine learning provides attackers and scammers with a powerful new tool that can probe for weaknesses, and make ransomware targeting more convincing and effective, among other dangers. But, used properly, these technologies also give defenders new resources that can make security technologies more effective and intuitive, while helping to ameliorate cyber workforce shortages.

I am currently the Chief Trust Officer of SentinelOne, a company that uses AI to help defend small to large enterprises, governments and nonprofits around the world. I am also a lecturer in the Computer Science and International Relations departments at Stanford University, where I teach classes in cybersecurity and online safety that include the creation of new AI tools by my students. I previously served as the Chief Information Security Officer at two large public companies, Facebook and Yahoo, and have consulted with hundreds of companies around the world both before and after serious cybersecurity incidents. I just finished a two-year term as a member of the DHS Cybersecurity Advisory Committee, am currently a member of the Aspen Institute U.S. Cybersecurity Working Group and also advise the NATO Cybersecurity Center of Excellence.

In my testimony, I will draw on my personal experience as a career cybersecurity professional to lay out a brief picture of the current security environment, with a focus on the ransomware threat, as well as some thoughts on how we can harness the power of AI in a safe way. I will also offer my thoughts on how we can build off of recent federal policy efforts like President Biden’s AI Executive Order<sup>2</sup> to create an effective and sustainable framework for the safe use of AI in the public and private sectors.

---

<sup>1</sup> [SentinelOne](#)

<sup>2</sup> [Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence | The White House](#)

## The Current Situation in the Field

Over the last two decades I have helped investigate and respond to dozens of attacks against American businesses. Before addressing how AI could impact cybersecurity I wanted to offer a handful of observations from the past year:

- **Cyber-extortion is a massive risk for companies of all sizes.** While we do continue to see interesting and important intrusions from state-sponsored actors, the baseline risk for every company in the United States, no matter their size or industry, are the professional extortion groups - cybercriminals.
- **Extortionists are getting bold and inventive.** Extortion groups are regularly demanding massive ransoms, in the range of \$40-60 million. When the victim (appropriately) attempts to negotiate this to a more reasonable level, threat actors use text messages to employees, emails to vendors and customers, ACH theft from the bank accounts of counterparties, and even the threat of Securities and Exchange Commission (SEC) investigation<sup>3</sup> to try to drive negotiations forward.
- **The current sanction regime has made paying more complicated but not less logical.** The cybercrime wave has created a niche industry of companies that specialize in tracking extortion groups. During several recent incidents, my clients were told by these specialists that, from the decision to pay the ransom being made, it would take five to seven days for the ransom payment to reach the threat actor. Most of that time is spent with sanctions compliance work, and given that these groups don't operate on layaway, the delay makes the strategy of paying to speed up recovery of systems less effective.
- **The SEC is creating new requirements that confuse cyber reporting.** In 2022, Congress passed the Cyber Incident Reporting for Critical Infrastructure Act<sup>4</sup> (CIRCIA) to standardize the process of reporting intrusions to the US Government. Via CIRCIA, Congress specifically directs CISA to be the focal point of cyber incident reporting in the US Government. The SEC has ignored Congress' will and imposed new reporting requirements for public companies that do not consider the difficult tradeoffs involved in public disclosure. While it is important that public companies are honest with investors, the requirement to file statements during the opening hours of a response and negotiation period gives the attackers more leverage and distracts from key response steps during the period when containment is almost never guaranteed. Threat actors have noticed and have used threats of SEC reporting to gain leverage, as previously referenced.
- **Many companies are vulnerable due to their traditional Microsoft architecture, and upgrading is extremely expensive.** Microsoft continues to dominate the enterprise information technology stack, with many organizations still running the same traditional on-premise Active Directory infrastructure that Microsoft recommended for years. Unfortunately, professional attackers have become extremely adept at finding and exploiting the common weaknesses in this kind of corporate network. More modern designs for Windows networks now exist, but generally require companies to subscribe to monthly Microsoft cloud services that many organizations find prohibitively expensive. The cost of Microsoft's licenses continue to slow down the adoption of modern technologies, and are also related to the forensic challenges faced by multiple Government agencies that struggled with

---

<sup>3</sup> [Ransomware group reports victim it breached to SEC regulators | Ars Technica](#)

<sup>4</sup> [Cyber Incident Reporting for Critical Infrastructure Act of 2022 \(CIRCIA\) | CISA](#)

investigating the breach of Microsoft's systems due to the lack of logging in their base cloud subscriptions<sup>5</sup>.

- **Legal risks bend companies away from smart, transparent responses.** The first call by a company during a breach is to outside counsel, and due to privilege concerns the cyber-lawyers are represented on every single call or email thread. I have worked with some excellent attorneys on breaches, but the over-legalization of executive decision-making is keeping companies from making smart, ethical, and transparent decisions because doing so might increase their risk of Department of Justice (DOJ), SEC or shareholder action in the future. I once worked a breach where there were four law firms on every call, representing various parties at the company, which did not engender long-term, transparent decisions from the executive team.
- **It has become very hard to hire qualified Chief Information Security Officers (CISOs).** There is a massive deficit of security leadership with the technical and leadership skills necessary to guide large enterprises through a cyber crisis. Recent actions by the SEC to lay the blame for systemic security failures on the CISO are exacerbating this problem<sup>6</sup>, and I personally know two well-qualified people who have passed up promotions to CISO roles due to the personal risk they would be taking.

## The Impact of AI on Cybersecurity<sup>7</sup>

I mention these facts because I expect that the AI revolution that we are just beginning to witness will have massive impacts on the struggle to secure US businesses from attacks, and that the basic roles played by security operators will look quite different in only a few years. This is mostly a good thing! As you can tell from my observations above, while great strides have been taken by Congress, the Executive Branch and companies across the nation, I am overall quite pessimistic about the current state of cybersecurity in the United States. One of the major drivers of our challenges is a lack of qualified individuals compared to the huge number of organizations that require them. While other industries rightfully fear AI replacing the jobs of humans, I am hopeful that the next several years will lead to AI developments that help close the massive gap in cybersecurity skills while leaving plenty of high-paying jobs for humans supervising AI agents.

Some of the benefits for defenders will include:

- Automated agents that can **sort through petabytes of security events and provide real-time visibility**<sup>8</sup> across a huge network. Our industry has done a great job of creating a huge amount of security telemetry from the tens of thousands of computers and other devices in a typical corporate network, but we have yet to put the ability to understand that data into the hands of your typical IT team.
- **AI-operated security operations centers (SOC)**, where the difficult 24x7 work of responding to security alerts will be left in the hands of computers while humans are woken up to provide oversight and to double-check the decisions of the AI agents. AI-enabled investigations will be much faster

---

<sup>5</sup> [Microsoft under fire after hacks of US State and Commerce departments | Reuters](#)

<sup>6</sup> [Cyber Chiefs Worry About Personal Liability as SEC Sues SolarWinds, Executive - WSJ](#)

<sup>7</sup> In this testimony I will restrict myself to discussing the impact of AI on the traditional information security field. I also have concerns around the impact AI could have on the manipulation of the US public by our foreign adversaries which I discussed in my testimony to the bipartisan Senate AI Forum. [Alex Stamos Statement - AI Insight Forum on Elections and Democracy](#)

<sup>8</sup> SentinelOne is one of several companies working to deploy LLMs and other AI models to this end: [Purple AI | Empowering Cybersecurity Analysts with AI-Driven Threat Hunting, Analysis & Response - SentinelOne](#)

and simpler for defenders, allowing them to make plain-English queries like “Show me all the computers that spoke to our secure network in the last eight hours” instead of struggling to get the exact syntax right on a search like:

```
ip.addr in {10.10.0.0 .. 10.10.0.254, 192.168.1.1..192.168.1.50}
```

- Real-time **analysis of unknown binaries, user behaviors and potentially malicious scripts** in a manner that most IT workers can understand. “Figure out what this potentially malicious piece of code does” used to be a question answered by a highly-skilled individual with a disassembler and debugger, and only the most highly resourced security teams can have such professionals as full-time staff. AI systems that can supplement these skill sets and provide plain-English explainability of complex programs will be hugely beneficial to defenders.
- More **flexible and intelligent response automation**. Many security coordination tools require a huge amount of effort to initially configure and are based upon fragile, human written rulesets. AI systems that respond to attacks in ways not fully foreseen by human defenders are both a scary idea and also likely necessary to cope with future attacks.
- Software development tools that point out insecure coding patterns to software developers in real-time, well before such bugs can make it into production systems. Reducing security flaws upstream is a much cheaper solution to our overall software security challenges than trying to patch bugs later.

It is also likely, however, that AI will be useful to attackers in several ways:

- AI could help attackers **sort through the billions of exposed services** they regularly scan to **automatically exploit** and install malware after the release of new flaws. This already happens, using human-written scripts, but AI could become a competitive advantage for groups that are able to use it to move faster and automate currently manual exploitation steps. Ultimately, speed kills in cyber, and AI may give attackers a new advantage.
- We will start to see regular **exploit creation via binary analysis**. Just as it requires specialized skills to analyze advanced malware it also requires specialized skills to write it, and there has already been research into using AI to create stable exploit code just through analyzing vulnerable programs with minimal human guidance.
- Smart malware that operates free of human direction or Command and Control (C2). AI could create new opportunities for criminal organizations to create **smart malware that operates behind air gaps**<sup>9</sup> or moves through networks intelligently, choosing the correct exploits and escalation paths without human intervention.
- Large Language Models are already **automating the work of social engineering** and ransom negotiations. Transformer tools are actively being used by cyber criminals to write more effective communications, including random demand emails, overcoming prior limitations in their grasp of the English language.

---

<sup>9</sup> The best example of malware with this capability is Stuxnet, which clearly required large amounts of intelligence around the design of the Natanz facility. Smart malware that does not require this kind of pre-existing knowledge is a goal of attackers and a nightmare for defenders.

It is quite possible that we are moving towards a world where the “hands on keyboard” actions currently performed by human attackers and defenders are fully automated, while small groups of experienced people supervise the AI agents that are automatically exploiting networks or fighting back against those exploits. Defenders may currently have an advantage in this space, as there has already been a decade of investment and research by security vendors into the defensive application of AI, however, we should not expect it to take long for attackers to catch up. That will be true for both the groups that hack for money and those who work for America’s adversaries.

### **The Near Term AI Policy Landscape**

President Biden’s AI Executive Order gave broad responsibilities to the Department of Homeland Security (DHS) and CISA, in particular, to aid the implementation of responsible, safe use of AI. The Order tasks CISA with developing guidance for critical infrastructure operators, and collaborating with public and private stakeholders to develop policies and practices around the use of AI<sup>10</sup>.

This is a critical mission, and just one of many that CISA has, and will continue to perform. The creation of a defense-only, non-regulatory agency that can support and partner with US companies was a great step by the 115th Congress and President Trump, and Congress should continue to ensure that CISA has the resources it needs to carry out this mission in an effective, responsive, and timely way. As cyber incident reporting requirements are built out pursuant to CIRCIA, Congress should continue to support CISA as the focal point for these reports, as well as response and remediation, and should work to de-conflict the various reporting requirements being invented by agencies outside of Congress’ direct recommendations.

As AI technologies evolve, it is important for policymakers to adopt nimble policies and safeguards made in careful collaboration with the private sector, and civil society groups representing a broad cross section of the country. As lawmakers carry out this vital but difficult mission, it is important that every effort is made to nurture and harness the positive benefits of AI, especially in the realm of security. Too many regulatory discussions around AI assume that only a handful of large American companies will dominate the space and can be utilized as chokepoints for preventing the malicious use of AI or spread of fundamental capabilities. This point of view is misguided and has led to warped regulatory priorities in the European Union and elsewhere.

The truth is that the AI genie is out of the bottle. There will be no reversing the spread of fundamental knowledge around modern AI techniques around the world. My Stanford students regularly use or even create new AI models as part of their classwork, and the amazing advances in open-source foundation models has demonstrated the capability of crowds of people to compete with US tech giants.

The spread of AI into every corner of personal and enterprise computing is inevitable. Congress should focus on encouraging responsible, thoughtful applications of these technologies and on maintaining the competitiveness of American champions instead of trying to control the spread of AI knowledge. America’s adversaries, and cyber criminals at home and abroad are sure to use these capabilities at every opportunity. It is critical that new regulations around the use of AI, however well intentioned, don’t hinder the ability of defenders to innovate and deploy these technologies in a beneficial way.

Thank you again for having me here today. I look forward to your questions.

---

<sup>10</sup> CISA’s initial output on this topic has been [published in tandem with the UK NCSC](#).