Harnessing Next-Generation AI Solutions for Enhanced Federal Security Operations

Highlights from a Roundtable hosted by the Advanced Technology Academic Research Center (ATARC) in partnership with SentinelOne, April 2024

Federal experts came together to discuss the challenges and opportunities of using generative AI to enhance Federal security. Panelists explored the multifaceted impacts of AI within cybersecurity, encompassing economic shifts, policy evolution, tech vulnerabilities, and the potential for developing global AI standards.

Opportunities and Challenges with AI in Cybersecurity

The introduction of generative AI into the public domain has brought both significant opportunities and challenges to national security. While panelists are encouraged by the positive outcomes of recent AI use cases, they also acknowledge the complexity of this moment.

Panelists note that certain industries, sectors, and federal departments are able to innovate and experiment with AI more so than others. For instance, the banking and finance sector lags in AI adoption due to stringent regulatory controls and significant risk aversion. Similarly, several Federal agencies have prohibited generative AI until more firm regulations are in place.

"The integration of artificial intelligence into our cybersecurity infrastructure presents not only vast opportunities, but also significant challenges that demand urgent attention and strategic action."

While these controls intend to mitigate risk, they are also inhibiting opportunities to experiment with AI – a requirement for innovation. The uneven implementation of AI could become problematic for some industries and departments, placing them at a significant disadvantage.

Impacts of Generative AI: Economy, Cybersecurity, and Policy

Panelists discussed the sprawling impacts of generative AI on the economy, cybersecurity, and policy development. AI's ability to automate routine tasks raises questions about workforce alignment and the potential impact on labor costs. As this technology shift plays out, panelists question the implications on future Federal budgets.

From a cybersecurity perspective, generative AI presents numerous challenges. AI's reliance on vast datasets creates friction with privacy regulations. While most data is aggregated, there is still the potential for insights and patterns to emerge despite privacy protections within specific documents and individual datasets.

Beyond data concerns, security vulnerabilities can be introduced through user prompting. When evaluating AI software, agencies should consider who has access to stored data, whether developer rights provide access to data, and whether investors can obtain information contained within prompts. These overlooked transparency gaps introduce the risk of data breaches and data misuse. Even with robust policies and secure systems, the threat of user error or insider threats remain significant.

The malicious use of AI to enhance social engineering schemes and vulnerability research is a growing concern, because it lowers the barrier to entry for cyber criminality. Roundtable participants also discussed challenges stemming from "shadow AI", where users deploy generative AI models outside established governance and perimeters. The lack of visibility in model interactions necessitates better monitoring tools, data classification standards, and internal data segmentation.

Considerations about AI in SaaS products

The integration of generative AI into SaaS products introduces another set of vulnerabilities and risks for government agencies. Roundtable participants agree that managing risk starts with clear terms or service and vendor agreements.

AI is often embedded into products without the government's knowledge or prior approval. Currently, the onus is on agencies to conduct due diligence and better understand the vendor's business processes that could introduce risk. Prioritizing SBOMs is an important step for agencies engaging in any new software, but especially AI software. The level of detail should include not only information about software parts, but also who owns the parts and who may have invested in the materials.

Without proper regulations, evaluating and approving AI SaaS solutions becomes onerous for individual agencies and inconsistent across government. However, becoming FedRAMP authorized is a notoriously slow and tedious process, yet without these authorizations, agencies cannot easily pilot new solutions, learn and refine prompt engineering, and adapt to rapidly evolving technology. While necessary, regulatory hurdles risk leaving the public sector lagging behind the private sector's innovation pace.

"The private industry is going around driving cars. We're still in horses and buggies."

Regulatory Challenges

Panelists discussed several regulatory and ethical challenges surrounding AI use. Because of AI's reliance on data, agencies must work to protect individual privacy and intellectual property rights while striving for innovation. This requires a well educated workforce knowledgeable about policies and ethical implications of AI. This is crucial to avoid misinterpretation of policy and the potential for over-enforcement.

"When I see people having problems with policy enforcement, it's because somebody is making a mistake, misreading the policy, over-enforcing a policy, or enforcing a policy where it doesn't belong…The knowledge base is not where it needs to be."

Panelists are predicting that smaller businesses without mature risk management practices and are embedding AI into products will soon face significant regulatory challenges. One panelist foresees existing regulatory agencies, such as the Federal Trade Commission, playing a role in AI oversight. Understanding AI decision-making and explainability during the contracting and procurement process will become paramount once AI regulation gets underway.

Compared to Europe's approach to AI regulation, the regulatory environment in the United States actually fosters a thriving innovation ecosystem. Instead of preventing regulatory challenges from occurring, the United States tends to correct bad outcomes once they occur.

Responsible AI Standards

Achieving a successful balance between responsible AI use and innovation necessitates transparency and accountability from stakeholders worldwide. Ideally, achieving transnational consistency in AI standards would streamline collaboration, drive alignment in AI development, and reduce the need to continually relearn rules.

In the warfighting arena, AI is changing the conflict landscape and requiring policy makers to develop new rules of engagement for warfighters. The challenge will be consistent interpretation and implementation of AI standards across boundaries. At the very least, Federal agencies need to agree on a standard definition of AI trustworthiness, so if the software meets one agency's requirements, another agency will know if it can meet theirs.

Final Thoughts

Agencies must work to incorporate responsible AI into policies in such a way as to not stifle creativity. This is especially important so the government can respond offensively to threats and keep a technological edge.

Ultimately, successful and responsible AI deployment hinges on empowered users who understand the systems, policies, and ethical considerations involved. It's unrealistic to assume anyone is receiving formal or informal education on responsible AI use, so the onus is on agencies to provide learning opportunities to its workers.