REQUEST FOR COMMENT RESPONSE

SAFE AND RESPONSIBLE AI IN AUSTRALIA: DISCUSSION PAPER

July 26, 2023

I. INTRODUCTION

In response to the Australian Government's Department of Industry, Science and Resources' ("Department") request for feedback on its *Safe and Responsible AI in Australia: Discussion Paper* ("Discussion Paper"), CrowdStrike offers the following views.

We approach these questions from the standpoint of a leading international, US-headquartered, cloud-native cybersecurity provider that defends globally distributed enterprises from globally distributed threats. CrowdStrike offers insights informed by multiple practice areas: cyber threat intelligence; proactive hunting, incident response and managed security services; and an AI-powered software-as-a-service cybersecurity platform and marketplace. Accordingly, this perspective is informed by CrowdStrike's role in protecting organizations from data breaches and a variety of other cyber threats.

II. COMMENTS

CrowdStrike welcomes the Department's ongoing efforts to ensure responsible innovation in Artificial Intelligence ("AI"), mitigate its potential risks and increase public trust and confidence in its development. We recognize that the Department has taken an important step to discover potential gaps in the existing domestic governance landscape and identify any possible additional AI governance mechanisms to support the development and adoption of AI.

As the Discussion Paper notes, AI is introducing new capabilities in many different sectors like medicine and engineering, and disrupting normal ways of doing business. The Discussion Paper correctly states that AI is evolving at a rapid pace and creating benefits and risks across many sectors and aspects of life. The cybersecurity sector is no different, with AI enhancing security capabilities while in other respects elevating new risks or threats. From a cybersecurity standpoint, the use of AI to detect threats is an enormous advantage. As adversaries continue to evolve and find new ways to target

victims,[1] organizations must increase their emphasis on cybersecurity practices that leverage the most effective technologies, including AI. AI is the best tool cyber defenders have to identify and prevent zero-day attacks and malware-free attacks. While the public discourse around AI has grown exponentially in the last year, AI in cybersecurity is not a new concept. CrowdStrike has deployed AI at scale across tens of millions of endpoints for prevention, dating back ten years.

While we do not have feedback on every question raised in the Discussion Paper, we do want to offer several points that may be of value to the Department on its journey to safe and responsible AI in Australia.

> **5. *Are there any governance measures being taken or considered by other countries (including any not discussed in this paper) that are relevant, adaptable and desirable for Australia?***

The Discussion Paper capably surveys international efforts on AI. CrowdStrike recently responded to a request for comment by the U.S. White House Office of Science and Technology Policy[2] ("OSTP"), which is developing a National AI Strategy.[3] That request for comment raised a few questions not included in the Department's Discussion Paper. OSTP asked about the national security benefits associated with AI; how AI can rapidly identify cyber vulnerabilities in existing critical infrastructure systems; and how it can be used to generate and maintain more secure software and hardware. While the governance measures regarding these complex issues are still being debated, proactively identifying areas for alignment, where possible, will create a coherent environment for continued innovation and make compliance more straightforward for companies operating internationally.

---

[1] Cyber threat actors are constantly evolving and finding new ways to target victims, including private sector entities and government agencies involved in national security. Illustrative of this, in CrowdStrike's 2023 *Global Threat Report*, we observed a notable surge in identity-based threats, cloud exploitations, and malware-free attacks. Adversaries continued to move beyond malware to gain initial access with malware-free activity accounting for 71% of all detections in 2022 (up from 62% in 2021). A contributing factor to this shift was the rate at which new vulnerabilities were disclosed and the speed with which adversaries were able to operationalize exploits. Further, we found a 95% increase in cloud exploitation by threat actors, over 30 new adversaries, and numerous new ways that eCrime actors weaponize and exploit vulnerabilities. *CrowdStrike Global Threat Report*, 2023. https://www.crowdstrike.com/global-threat-report/

[2] Please contact policy@crowdstrike.com for a copy of our comments.

[3] White House Office of Science and Technology Policy (OSTP) "*Request for Information National Priorities for Artificial Intelligence*" *May* 23, 2023. https://www.whitehouse.gov/wp-content/uploads/2023/05/OSTP-Request-for-Information-National-Priorities-for-Artificial-Intelligence.pdf

9. Given the importance of transparency across the AI lifecycle, please share your thoughts on where and when transparency will be most critical and valuable to mitigate potential AI risks and to improve public trust and confidence in AI?

In our view, many AI safeguards, including specific transparency or interpretability requirements, should account for the use case. At a minimum, we suggest a distinction between "Consumer-facing AI" versus "Enterprise AI" applications. Consumer-facing AI applications may be more likely to involve impacts to specific users, people, or groups, or have other social implications. Enterprise (B2B) technologies that use AI may improve business processes and drive efficiency, like optimizing a maintenance schedule or reducing energy consumption. Different safeguards may be appropriate in each instance, but should relate specifically to potential risks.

Consumer-facing AI use cases may benefit from having explainability requirements, where the AI provider is required to be able to account for the knowledge base or data set the AI tool is drawing from (e.g. quality, accuracy, amount of data going into the algorithm). In contrast, such requirements should not be applied to enterprise solutions that leverage AI. Blanket requirements for algorithmic explainability could have negative impacts on innovation and copyright, and would lack appropriate context. Therefore, in enterprise use cases, a contract often clarifies permissible uses of data for specific AI purposes.

11. What initiatives or government action can increase public trust in AI deployment to encourage more people to use AI?

The widening adoption of AI periodically raises concerns about automated decision-making, surveillance, algorithmic bias, and other risks or negative externalities. But in the discourse on AI, policy makers and the public must also consider that AI also has the opportunity to drive positive social outcomes; is already widely deployed in important instances driving such outcomes; and creates the opportunity for innovation in a variety of important sectors, including industries such as cybersecurity, medicine, and education.

As discussed in our general comments above, the use of AI to detect cyber threats provides a significant benefit.[4] Considering that cyberthreats can have potentially

[4] Michael Sentonas, *How Artificial Intelligence is Becoming a Key Weapon in the Cybersecurity War*, CrowdStrike Blog, October 24, 2017. https://www.crowdstrike.com/blog/how-artificial-intelligence-is-becoming-a-key-weapon-in-the-cybersecurity-war/.

disastrous consequences for our democratic institutions, infrastructure, and economies, the government could convey the positive message that AI is helping to protect citizens from these threats. We would be happy to work with stakeholders to identify case studies to these ends.

With respect to encouraging the use of AI, the advent and rapid productization of LLMs has opened a new frontier of experimentation. While many early uses are consumer-facing applications focused on creative or artistic endeavors, we see enormous opportunities within enterprise software applications. For example, CrowdStrike leverages LLMs to assist analyst workflows and to make other security analyst tasks more efficient. This capability (coined "*Charlotte*") utilizes CrowdStrike's highest-fidelity security data, which includes the trillions of security events captured in the CrowdStrike Threat Graph, asset telemetry from across users, devices, identities, cloud workloads, and threat intelligence. The use of this knowledge base drives efficacy, actionability, and relevance, while also addressing the risk of "hallucination." Further, the natural language interface seeks to make cybersecurity responsibilities more broadly accessible. Our goal with *Charlotte* is to help close the cybersecurity skills gap and improve the response time so users can stay ahead of adversaries – boosting security across organizations. Today, we see this use of LLMs as one of the most relevant to improving security outcomes in the near- to mid-term.

14. Do you support a risk-based approach for addressing potential AI risks? If not, is there a better approach?

We support the use of a risk-based approach to assessing the overall need, scope, strength, timing, and sequence for AI-related law(s), policy(ies), and regulation(s). Resulting measures, however, should in most cases be principles-based. Outcome-based measures can be appropriate in some contexts, but this approach can be challenging with frontier technologies or areas with adaptive adversaries. Rules- and standards-based measures should be used sparingly, if at all. Selecting the right approach for a given area is complex, and we recommend additional consultation with stakeholders through iterative comment requests.

More broadly, our view is that for AI, like for any other technology, the context in which it is used, rather than the mere fact that it is incorporated, is material. Consequently, regulating AI for the sake of the technology rather than its application is not the best approach to foster-innovative solutions to difficult problems. CrowdStrike recommends that policy addressing AI should be targeted and specific. Sound policy

should address potential threats and risks, and ultimately support innovation by clarifying research constraints and related legal and contractual issues. (See also our response to Question 18, below.)

> ***18. How can an AI risk-based approach be incorporated into existing assessment frameworks (like privacy) or risk management processes to streamline and reduce potential duplication?***

An AI-risk based approach can be incorporated into existing assessment frameworks (like privacy) to streamline and reduce potential duplication. Specifically on privacy and AI, CrowdStrike recently provided feedback in response to the Attorney General's requests for comments on the Privacy Act.[5] In our comments, we noted our opposition to Proposals 19.1, 19.2, and 19.3, which suggest policies that focus on protecting individual rights by giving them the right to object to how a particular technology (AI) uses their information.

Instead, as we noted above in our response to Question 14, for AI, like for any other technology, the context in which it is used, rather than the mere fact that it is incorporated, is material. Consequently, relying upon a right to object to a particular technology or data processing methodology is not the best approach to protect privacy rights in an ever-evolving technological landscape. Instead, we recommended protecting the rights of Australians through a technology-neutral approach. When creating regulations on the safe use of AI, we suggested that Australia should consider adopting language similar to the General Data Protection Regulation's ("GDPR") requirement that organizations implement safeguards "appropriate" to the risk to protect personal information. This approach incentivizes organizations to take into account modern, rapidly-evolving data breach risks posed by cybersecurity threats from e-crime, 'hacktivist', and nation state actors using tactics such as ransomware, supply chain attacks, or malware-less intrusions.

The same principles of data privacy apply across various uses, including AI. Data protection involves integrity, confidentiality, and availability. Accordingly, the context of data processing activities will decide the best ways for protection.

III. CONCLUSION

[5] CrowdStrike Request for Comment Response, "AUSTRALIA PRIVACY ACT REVIEW REPORT 2022," March 31, 2023. https://www.crowdstrike.com/wp-content/uploads/2023/06/australia-privacy-act-review-report.pdf

The Department's Discussion Paper provides a thoughtful analysis of a complex legal and policy area. As the Discussion Paper moves forward, we recommend continued engagement with stakeholders. Finally, because the underlying technologies evolve faster than law and policy, we recommend that any updates focus on principles rather than prescriptive requirements and include a mechanism for periodic revisions.

IV. ABOUT CROWDSTRIKE

CrowdStrike®, a global cybersecurity leader, is redefining security for the cloud era with an endpoint protection platform built from the ground up to stop breaches. The CrowdStrike Falcon® platform's single lightweight-agent architecture leverages cloud-scale AI and offers real-time protection and visibility across the enterprise, preventing attacks on endpoints on or off the network. Powered by the proprietary CrowdStrike Threat Graph®, CrowdStrike Falcon correlates over 3 trillion endpoint-related events per week in real time from across the globe, fueling one of the world's most advanced data platforms for security.

With CrowdStrike, customers benefit from better protection, better performance and immediate time-to-value delivered by the cloud-native Falcon platform.

There's only one thing to remember about CrowdStrike: We stop breaches. Learn more: https://www.crowdstrike.com/.

CONTACT

We would welcome the opportunity to discuss these matters in more detail. Public policy inquiries should be made to:

Drew Bagley CIPP/E
VP & Counsel, Privacy and Cyber Policy

Karen Kaya
Senior Manager, Public Policy

Email: policy@crowdstrike.com

CrowdStrike owns other trademarks and service marks, and may use the brands of third parties to identify their products and services.

***