5 September 2025

Horizon 2 Cyber Strategy Consultation
Department of Home Affairs

BY EMAIL: CSSH2@homeaffairs.gov.au

Google welcomes the opportunity to provide feedback on the discussion paper *'Charting New Horizons - Developing Horizon 2 of the 2023 - 2030 Australian Cyber Security Strategy'*.

Google supports the Australian Government's goal of uplifting Australia's cyber security posture across the economy, and is committed to working with the Australian Government, States and Territories, and across industry as a trusted partner.

To make this vision a reality, we recommend that Horizon 2 builds on the existing strategy to continue advancing technology policies that put user protections first, promote innovation and technology modernisation, and explore new models of public-private partnership for collective defence.

The digital landscape is rapidly evolving, presenting both unprecedented opportunities and significant challenges that necessitate a collaborative, multi-stakeholder approach between government, industry, and civil society. As a global leader in cybersecurity, Google is committed to partnering to combat threats and improve security across the digital ecosystem for everyone.

Google has a long history in building secure infrastructure and helping to create and define cybersecurity best practices. Our approach to security – which emphasises a **focus on the fundamentals** of secure software development; an **open and transparent approach** to improving the ecosystem; and **pushing the frontiers of advanced security technologies** – have served us well and can serve as a useful roadmap for policymakers.

We protect our users and enterprise customers by providing industry-leading security. We are committed to keeping users and customers, and Internet infrastructure more broadly, secure. We do this in part by contributing to international security standards, sharing best practices, templates, developer tools, and providing other integrated solutions that make security stronger and easier to implement. And of course by offering secure services to our customers and users and implementing a shared fate approach to risk management rather than delineate where our responsibility ends.

Australia's Horizon 2 Strategy and Beyond

Supporting cybersecurity uplift for business

Google recognises that SMBs face unique challenges in cybersecurity, often lacking the resources and expertise of larger enterprises. Our approach is to make advanced security accessible and effective for businesses of all sizes:

- **Built-in Product Protections:** Many Google products, widely used by SMBs, come with inherent security features. Gmail automatically blocks over 99.9% of spam, phishing, and malware. Chrome's Safe Browsing and Password Manager proactively protect users from dangerous websites and credential theft. Google Play Protect scans billions of installed apps daily, blocking malware installations.
- **Google Cloud Offerings:** We provide scalable and cost-effective security solutions through Google Cloud that are particularly beneficial for SMBs. For instance, *reCAPTCHA Enterprise* helps protect websites and mobile applications from fraud, bots, account takeovers, and credential stuffing. This service protects over 6 million websites from fraud and bots.
- **Risk Protection Program:** This innovative program, in partnership with leading cyber insurers like Beazley, Chubb, and Munich Re, helps Google Cloud customers, including SMBs, assess and manage their security risks. The *Cyber Insurance Hub* scans Google Cloud workloads, provides security recommendations, and generates reports mapping configurations against industry standards (CIS Benchmarks). These reports can be shared directly with partner insurers for streamlined underwriting and potentially more competitive pricing.
- **Education and Workforce Development:** To address the cybersecurity talent gap that disproportionately affects smaller entities, Google launched the Google Cybersecurity Certificate. It trains individuals for entry-level cybersecurity roles, equipping them with practical skills in Python, Linux, SQL, and SIEM tools like Chronicle and Splunk.
- **IT Modernisation:** Cloud adoption provides SMBs with a fast and cost-effective way to implement robust security programs that were previously only available to well-resourced companies. This allows them to benefit from continuously updated defences without significant internal investment in security infrastructure.

We regularly publish information on cyber threat trends that we see impacting SMBs and guidance on what they can do to enhance their cybersecurity posture. This includes a continued focus on cybersecurity best practices, such as strong passwords, data backups, and response and recovery plans, ensuring cybersecurity is a resourced priority and encouraging digital modernisation.

Building Safe technology

"Secure by Design" is a fundamental principle woven into the fabric of Google's product development and operations. It means security is considered and integrated at every stage of the product development lifecycle, from initial design to deployment and ongoing maintenance, rather than being "bolted on" as an afterthought.

Key aspects of our approach include:

- **Foundational Security:** We prioritise secure software development practices, embed security controls like multi-factor authentication, and implement zero-trust architecture by default. This approach minimises vulnerabilities from the outset.
- **Proactive Threat Elimination:** Our focus is on addressing the root causes of vulnerabilities to eliminate entire classes of threats. This mindset has evolved since incidents like Operation Aurora in 2009, leading to a continuous refinement process embedded in our culture, products, and services.
- **Secure AI Framework (SAIF):** Recognising the emerging risks from AI, Google developed SAIF,a conceptual framework to guide the secure development and deployment of AI systems. SAIF advocates for applying decades of software security learnings to AI, ensuring systems are secure-by-design and default, and addressing new vulnerabilities like prompt injection attacks. We actively work with partners and governments to drive adoption of SAIF and safeguard privacy, mitigate unexpected actions, and ensure human oversight and transparency in AI agents. See further detail below.
- **Continuous Updates and Defences:** Our cloud platform is continuously updated with the latest defences, ensuring that customers benefit from cutting-edge security measures without manual intervention.

Data access and transfer across the economy

Numerous government regulations and laws require geographical locking of data as a form of security. Google's view is this creates complacency that data is secure when kept onshore in Australia, yet hackers are not geographically bound and will exploit data globally. Additionally, the Australian Government is seeking to make use of global computing power and innovation in AI to advance its standing, but this is at odds with government security policy and regulation. Google suggests that changing the view of sovereignty from geo-boundaries to sovereign control of your cloud services through encryption, sovereign key management, and strong access controls would provide the same, if not better, security over your data. This is achieved using industry leading encryption in all stages of data use, backed by the global security monitoring of Google, allowing rapid technology adoption and innovation promoting a technically advanced Australia. Additionally, resilience in the hosting of data is increased by not geo-bounding your data location should conflict or other events disconnect your data from your users.

Secure AI Framework
The potential of AI, especially generative AI, is immense. As innovation moves forward, the industry needs security standards for building and deploying AI responsibly. Google has introduced the Secure AI Framework (SAIF), a conceptual framework to secure AI systems. SAIF is designed to address top-of-mind concerns for security professionals, such as AI/ML model risk management, security and privacy – helping to ensure that when AI models are implemented, they are secure by default. The six core elements of SAIF are:

- Expand strong security foundations to the AI ecosystem
- Extend detection and response to bring AI into an organisation's threat universe
- Automate defences to keep pace with existing and new threats
- Harmonise platform-level controls to ensure consistent security across the organisation
- Adapt controls to adjust mitigations and create faster feedback loops for AI deployment
- Contextualise AI system risks in surrounding business processes.

Addressing concentration risk
A recent U.S. Cyber Safety Review Board (CSRB) report and similar observations underscore the significant risks associated with government over-reliance on a single technology vendor. Such "concentration risk" creates systemic weaknesses that can be exploited by adversaries, undermining an entire ecosystem if that vendor is compromised. To mitigate this, Google strongly recommends:

1. **Vendor Diversification:** Governments should actively adopt a multi-vendor strategy for their IT infrastructure and security tooling. This reduces the impact of a breach in any single vendor's products and enhances overall resilience.
2. **Procurement of Secure-by-Design Systems:** Security must be a non-negotiable criterion from the earliest design phases of products and throughout their lifecycle. Procurement processes should prioritise technologies built with inherent security rather than those that require security "add-ons".
3. **Integrating Security into Procurement Decisions:** Security assessments should be ongoing and consider a vendor's history of breaches, rather than ending with initial accreditation. Procurement officials should factor in past performance regarding security, alongside traditional metrics like cost and on-time delivery.
4. **Promoting Interoperability and Open Standards:** Governments should develop and promote open standards to ensure different systems and products can work together seamlessly. This facilitates easier replacement of insecure products and fosters a more competitive and diverse supplier ecosystem, benefiting innovation and mitigating concentration risk.
5. **Scrutinising Restrictive Licensing Practices:** Policies that impede a diverse supplier ecosystem and disincentivise innovation and customer choice should be investigated

and reformed. Anti-competitive lock-in practices ultimately limit choice and can compromise security.

6. **Accelerated Adoption of Modern Cloud Solutions:** Migrating from legacy, on-premises systems to modern cloud environments offers significant security advantages, including built-in protections, continuous updates, and the ability to scale defences rapidly. This shift also helps diversify infrastructure beyond potentially vulnerable legacy setups.

World-class threat sharing and blocking

Google uses our unique capabilities to detect and disrupt threat actors on our platforms. More actors than ever are engaged in cyber espionage and other attacks against organisations and individuals. At Google, we unify threat intelligence from multiple sources to provide unmatched visibility on malicious activity and threat trends.

On October 10, 2023, Google Cloud mitigated one of the largest DDoS attacks to date. In 2024, Mandiant spent 450,000+ hours investigating incidents. Google keeps more people and organisations safe online than anyone else in the world — billions of users and devices around the globe, blocking 100 million phishing attempts and classifying 10 billion websites and files every day. Over 7 million sites today use reCAPTCHA, our fraud and bot management solution that helps to stop attacks like credential stuffing and account takeovers.

We also make threat reporting and threat trend information and technical mitigation recommendations publicly available via our Google Cloud Threat Intel blog to enable network defence efforts broadly.

Threat Intelligence Sharing

Google recommends that governments prioritise establishing a cohesive, automated threat intelligence sharing ecosystem. While public-private cooperation has improved, it remains largely manual, slow, and constrained by restrictions. The goal is to move towards an operationalised model where intelligence enables real-time, automated blocking of threats at scale. To achieve this, governments should deepen public-private partnerships to ensure a consistent, bidirectional flow of actionable threat intelligence, as both sectors hold unique and crucial insights. They should also incentivise the use of open-source and supply chain security frameworks to facilitate automated, machine-to-machine exchange of threat indicators. Finally, conducting joint cyber exercises is essential to clarify the roles and responsibilities of government and industry, which will improve preparedness and ensure a swift, coordinated response in a crisis.

Google supports the implementation of a vulnerability disclosure program (VDP) in Australia. Such a program should aim to create a “safe harbour” for security researchers, protecting them from legal consequences for responsibly reporting vulnerabilities. This provides a crucial mechanism for organisations to identify and fix security flaws before they can be exploited by malicious actors.

Vulnerabilities should be coordinated through industry-recognised coordinated vulnerability disclosure (CVD) processes and not reported directly to governments. This approach ensures that the entity responsible for the vulnerable product or service is the first to be notified, enabling a direct and efficient path to remediation. The government's role should be to facilitate this process by providing a legal framework that encourages and protects good-faith reporting, rather than acting as a clearinghouse for vulnerability information. By incentivising businesses to adopt their own vulnerability disclosure policies, the government can leverage the global security research community to strengthen the nation's cyber defences.

Information Sharing and Analysis Centres

Google is a strong supporter of Information Sharing and Analysis Centres (ISACs) and has supported ISACs in Australia, such as the CI-ISAC. Information sharing is a vital tool in cybersecurity defence, and it has come an incredibly long way over the past several decades with security intelligence feeds from dozens of organisations (including Google Cloud) used to enhance cyber-defences and stop threat actors before they succeed.

Applying AI and end-to-end automation to support the ability to act on threat intelligence is crucial to accelerate response time, as is merging threat intelligence with organisational risk context. Interpersonal and cross-organisational communications are also hugely important. Some of that communications work involves framing security strategies properly so that business executives and boards of directors are actively involved, understand the crucial role they play, to ensure that security operations are prioritised, funded, and optimised.

From the network defence perspective, operators who can effectively share threat intelligence raise the cost to malicious actors of a successful attack. ISACs can help cut through the noise to give their member organisations clear signals. Their highly-trusted, peer-to-peer networks are very effective at sharing threat intelligence — including attackers' goals, objectives, tools, tactics, and procedures with immediate relevance to their community — and responding to that shared intelligence faster and faster.

Emerging technology and cyber defence

AI provides an unprecedented opportunity for building a new era of innovation. We can use these new tools to grow the economy, create jobs, accelerate scientific advances and give the advantage back to security defenders. When it comes to security opportunities — we're thrilled to be driving progress in three key areas: agentic capabilities, next-gen security model and platform advances, and public-private partnerships focused on putting these tools to work.

Last year, we announced Big Sleep, an AI agent developed by Google DeepMind and Google Project Zero, that actively searches and finds unknown security vulnerabilities in software. By November 2024, Big Sleep was able to find its first real-world security vulnerability, showing the immense potential of AI to plug security holes before they impact users.

Since then, Big Sleep has continued to discover multiple real-world vulnerabilities, exceeding our expectations and accelerating AI-powered vulnerability research. Most recently, based on intel from Google Threat Intelligence Group, the Big Sleep agent discovered an SQLite vulnerability (CVE-2025-6965) — a critical security flaw, and one that was known only to threat actors and was at risk of being exploited. Through the combination of threat intelligence and Big Sleep, Google was able to actually predict that a vulnerability was imminently going to be used and we were able to cut it off beforehand. We believe this is the first time an AI agent has been used to directly foil efforts to exploit a vulnerability in the wild. These AI advances don't just help secure Google's products. Big Sleep is also being deployed to help improve the security of widely used open-source projects.

These cybersecurity agents are a game changer, freeing up security teams to focus on high-complexity threats, dramatically scaling their impact and reach. In our recent white paper on AI agent security, we outline our approach to building AI agents in ways that safeguard privacy, mitigate the risks of rogue actions, and ensure the agents operate with the benefit of human oversight and transparency. When deployed according to secure-by-design principles, agents can give defenders an edge like no other tool that came before them. Collaboration across industry and with public sector partners is essential to cybersecurity success. We have always believed in AI's potential to make the world safer, but over the last year we have seen real leaps in its capabilities, with new tools redefining what lasting and durable cybersecurity can look like. Recent advances in AI have the potential to be game-changing, but what we do next matters. By building these tools the right way, applying them in new ways and working together with industry and Governments to deploy at scale, we can usher in a digital future that's not only more prosperous, but also more secure. We also continue to publish materials on how we're thinking about cybersecurity issues, and what businesses can do to uplift their posture.

Post-quantum cryptography adoption is rapidly becoming a reality, and the need for active deployment and implementation is becoming increasingly urgent. At Google, we're taking proactive steps to secure your data against the potential threats from cryptographically relevant quantum computers (CRQC). Our commitment is to safeguard your sensitive information, not just for today, but for the years ahead. We're also dedicated to sharing our knowledge and expertise, advising both our customers and the wider industry on best practices for navigating the transition to a quantum-safe future. As part of our years-long efforts to prepare for the arrival of a CRQC world, earlier this year we introduced quantum-safe digital signatures as part of Google Cloud Key Management Service. This follows strategy and development efforts going back a decade, including PQC tests in Google Chrome in 2016, protecting internal communications with PQC since 2022, and additional quantum-safe protective measures across our products and services.

Harmonising and simplifying regulation

Regulating cybersecurity at a national scale is complex, poses unique challenges, and carries high stakes. Regulatory and compliance regimes impact critical infrastructure resilience, economic development, the pace of innovation, and the daily lives of citizens. As a result, cybersecurity regulation should be carefully balanced: promoting strong cybersecurity baseline standards while allowing flexibility to account for evolving technology and the ever-changing threat landscape.

Google recommends a regulatory approach that is agile and focuses on aligning baseline requirements across sectors. The approach should also allow for additional sector-specific requirements that are complementary to, and not duplicative of, or in conflict with, those standard baselines. Regulatory agility will help reduce compliance burdens, enhance coordination, build public trust, and allow for a more resilient approach as threats change, new economic sectors emerge, and agency responsibilities change and shift over time.

Regulations must prioritise tangible outcomes over mere checklist compliance. Google's commitment to openness, interoperability, transparency, responsibility, a secure-by-design approach, intelligent security systems, and collaborative efforts can only be fully realised within such an adaptable regulatory environment. We urge governments to consider modernising cybersecurity regulations and creating a stable baseline based on existing, broadly accepted industry standards, that existing sectors can adhere to and future sectors can adopt as a reliable guide for improving security and resilience.

To achieve regulatory harmonisation, Google offers a few central recommendations:

- Leverage well-established standards and processes for any contemplated security baseline approach.
- Any harmonised standards should implement a risk-based approach - ensuring compliance options are aligned to specific risk levels or categories to maximise flexibility and efficiency commensurate with the level of risk associated with a particular technology, application, or use case.
- We encourage continued focus and clarity on reciprocity for different certification regimes wherever possible.
- Changing the view of sovereignty from geo-boundaries to sovereign control of your cloud services through encryption, sovereign key management and strong access controls.
- Harmonising international cybersecurity regulations, though a challenge, will offer significant rewards.

Organisations struggle with regulatory uncertainty, complexity, and fragmentation. We encourage consideration of a global harmonised approach to ensure enterprises and service providers can focus on security outcomes as a top priority. Google remains committed to the security of the digital ecosystem and would be pleased to consult on future cybersecurity regulations, and adjustments to existing requirements should such be under consideration.

Building a skilled workforce

The cybersecurity skills shortage is a critical concern for policy makers. Government agencies are increasingly vulnerable to cyber attacks, which threatens the integrity of essential public services and national security. The world has an urgent need for a more skilled cyber workforce.

Drawing on our experience, we believe a multi-faceted approach is essential to address the cybersecurity workforce shortage. This involves expanding career pathways to include individuals regardless of their background or prior experience. For example, our Google Cybersecurity Certificate, which is built and taught by Google experts, provides a flexible, online training program that prepares individuals for entry-level jobs in as little as three months. We are also committed to creating more hands-on learning opportunities. Through our Cyber Clinics program, we're providing grants and resources to universities to establish clinics where students can gain real-world experience by helping local organisations and critical infrastructure improve their cyber defences, creating a dual benefit of upskilling new talent while also strengthening community security. This year, we are investing $5M more to establish 11 cyber clinics – including in Australia – through our partnership with the Asia Foundation.

We welcome the opportunity to discuss our experience and to engage with Home Affairs as it undertakes consultation on the Horizon 2 Strategy.