Horizon

- Aus Gov
- 29/08/2025

List of questions to consider when making a submission to the Horizon 2 Public Discussion Paper

2.1 Outlook for Horizon 2

1. What trends or technology developments will shape the outlook over the next few years and what other strategic factors should Government be exploring for cyber security under Horizon 2?

Key Trends Shaping the Outlook

The cyber landscape of the next few years will be defined by three dominant and interconnected trends:

1. **The Hyper-Proliferation of the Attack Surface:** The explosion of connected devices (IoT/OT) and the shift to a permanent hybrid "everywhere enterprise" has dissolved the traditional network perimeter. This renders legacy prevention-focused security models insufficient, as every device, user, and cloud service becomes a potential entry point for adversaries.
2. **The Industrialisation and Weaponisation of AI:** Threat actors are moving beyond simple automation to leverage AI for creating hyper-realistic phishing, adaptive malware, and machine-speed attacks. This asymmetry, where attackers use AI against human-led defence teams, is unsustainable. The only viable long-term defence against AI-powered attacks is an AI-powered defence.
3. **The Amplification of Systemic Supply Chain Risk:** As digital ecosystems become more integrated, the compromise of a single managed service provider or software component can create a cascading failure across thousands of Australian organisations. This "single point of failure" risk is a prime target for adversaries seeking to maximise impact and highlights that organisational security is inextricably linked to the security of its partners.

Strategic Factors for Government Exploration

Given these trends, achieving the Horizon 2 vision requires moving beyond foundational standards and exploring more dynamic and forward-leaning strategic initiatives. We urge the Government to consider the following factors:

1. Champion a Nationwide Shift from Compliance to True Cyber Resilience: Foundational standards like the Essential Eight are vital, but they represent the starting line, not the finish line. A compliance-only mindset fosters a "checklist" approach to security that is ill-equipped for dynamic, sophisticated threats. The Government should strategically lead a national mindset shift towards **Threat Detection, Investigation, and Response (TDIR)**. This involves:

Evolving Policy: *Augmenting regulations to require not just preventative controls, but also demonstrated capabilities to detect, investigate, and respond to breaches in a timely manner, potentially by referencing frameworks like MITRE ATT&CK.*

Incentivising Capability Uplift: Creating programs, particularly for SMEs, that incentivise the adoption of modern Security Information and Event Management (SIEM) and analytics platforms. This would help businesses move beyond asking "Are we compliant?" to asking "How quickly can we detect and stop an attack?".

2. Foster an AI-Powered National Defensive Ecosystem: To counter AI-driven threats, Australia must cultivate a defensive ecosystem that leverages the same technologies. The Government can act as a catalyst by:

Leading by Example: *Making Federal Government agencies exemplars in the use of AI and behavioural analytics within their own Security Operations Centres (SOCs) to automate threat detection and reduce analyst fatigue.*

Building a Future-Ready Workforce: Partnering with industry and academia to develop a cyber workforce skilled not just in network engineering, but in data science, security analytics, and AI operations, which are the skills needed to manage the next generation of security tools.

3. Promote Radical Transparency and Collaboration in the Digital Supply Chain: The interconnected nature of our economy is a strength that can become a critical vulnerability. The Government can mitigate this systemic risk by fostering a culture of "collective defence."

Normalising Information Sharing: *Developing robust, real-time, and anonymised threat intelligence sharing platforms that extend beyond critical infrastructure to all sectors of the economy. This includes creating "safe harbour" legal protections for organisations that share breach information in good faith, removing the fear of punitive action and encouraging rapid, widespread alerts.*

By focusing on these strategic factors, the Government can build upon its foundational work, fostering a cyber ecosystem that is not just compliant, but genuinely resilient, intelligent, and collaborative enough to meet the challenges of 2028 and beyond.

• 2.2 Collaborating across all levels of Australian Government

2. Are there initiatives or programs led by State or Territory governments you would like to see expanded or replicated across other levels of government?

• 2.3 Monitoring progress in a changing world – a conceptual framework for evaluating cyber security outcomes

3. Does the high-level Model resonate and do you have any suggestions for its refinement?

While the high-level model provides an excellent framework, we offer the following suggestions for its refinement to enhance its precision, dynamism, and real-world impact, ensuring it measures not just activity, but true resilience.

1. Integrate Leading Indicators of Resilience, Not Just Lagging Indicators of Failure: The current model appears to focus on tracking outcomes, which are often lagging indicators (e.g., "cost of cybercrime," "number of incidents"). While important, these metrics tell you about past failures. To be truly responsive, the model must also measure leading indicators that signal an organisation's *capacity* to withstand and respond to an attack. We strongly recommend incorporating metrics central to Threat Detection, Investigation, and Response (TDIR), such as:

- **Mean Time to Detect (MTTD):** How long does it take for an organisation to become aware of a malicious intrusion? A decreasing national MTTD is a powerful indicator that visibility and detection capabilities are improving across the economy.
- **Mean Time to Respond (MTTR):** Once detected, how long does it take to contain and neutralise a threat? This metric directly measures the effectiveness of incident response processes, automation, and human expertise. Tracking MTTD and MTTR provides a real-time pulse check on our collective defensive capabilities, which is a more predictive measure of resilience than simply counting the number of successful attacks.

2. Introduce the Concept of "Fidelity of Interventions": The model's arrows represent "causal hypotheses." This causality can be strengthened by measuring the *quality* and *adoption depth* of an intervention, not just its deployment. For example, if an intervention is to "promote adoption of the Essential Eight," the model shouldn't just track how many businesses have adopted it (a binary measure). It should seek to measure *how well* they have adopted it. A low-fidelity implementation provides a false sense of security. The model could be refined to include maturity scores or fidelity ratings for key interventions, providing a more nuanced understanding of whether a given policy is having its intended effect or is merely being treated as a compliance checkbox.

3. Develop Metrics for Security Operations and Workforce Effectiveness: A critical, yet often overlooked, component of cyber resilience is the human element. The effectiveness of a nation's cyber defence is directly tied to the efficiency and expertise of its security professionals. The model should include metrics that track the health and capacity of our Security Operations Centres (SOCs) and cyber workforce. This could include:

- **Analyst Alert Fatigue:** Tracking the ratio of alerts to actual investigated incidents. A high ratio can be an early warning sign of workforce burnout and inefficient tooling, indicating that an intervention is needed to improve security analytics and automation.

- **Threat Intelligence Utilisation:** Measuring the speed and breadth of which actionable threat intelligence (e.g., from the ASD) is consumed and operationalised within organisations' defensive tools.

By integrating these refinements—a focus on leading resilience metrics like MTTD/MTTR, a nuanced view of intervention fidelity, and metrics for the human-led operational response—the Cyber Security Policy Evaluation Model can evolve from a powerful conceptual framework into a dynamic, real-time navigation system. This will provide a more accurate picture of the nation's cyber posture and be crucial for steering Australia towards its 2028 "North Stars" with confidence and agility.

4. Can you suggest any existing or new ways to collect data and feedback to monitor these outcomes?

Existing channels provide a valuable baseline, but their depth and focus can be significantly improved to feed the model with more meaningful metrics.

1. **Enriching National Surveys:**
    - **Current State:** The Government likely conducts periodic surveys of businesses and citizens on cyber security matters. These are often focused on high-level security postures and incident numbers.
    - **Suggested Enhancement:** Evolve these surveys to capture more sophisticated, quantitative resilience metrics. Instead of only asking, *"Do you have an incident response plan?"* (a yes/no question), the survey should include questions that probe capability, such as: *"What was your organisation's average **Mean Time to Detect (MTTD)** a critical security incident over the last 12 months?"* and *"What was your **Mean Time to Respond (MTTR)**?"* For SMEs, questions could be tailored to the adoption of specific, accessible security services. This shifts the focus from checking boxes to measuring actual operational readiness.
2. **Systematising Breach Reporting Data:**
    - **Current State:** Mandatory breach reporting to the OAIC and voluntary reporting to the ASD provide a wealth of data on incidents that have already occurred (lagging indicators).
    - **Suggested Enhancement:** Standardise and enrich the data collected during these reports. Beyond the nature of the breach, the reporting framework should be updated to require (where possible) structured data fields for the initial attack vector (aligned with

MITRE ATT&CK), the time to detection, and the time to containment. This transforms each breach report from a simple notification into a rich data set that directly feeds the MTTD and MTTR "North Star" metrics and helps identify systemic weaknesses across the economy.

3. **Deepening Public-Private Threat Intelligence Sharing:**
    - **Current State:** The ASD's Cyber Security Partnership Program is a strong foundation for threat intelligence sharing.
    - **Suggested Enhancement:** Invest in platforms that not only share threat signatures but also allow for the contribution of anonymised telemetry and metadata about defensive performance. The Government can act as a trusted intermediary, collecting anonymised data on what kinds of attacks are being blocked, what is being detected, and how long it takes, then sharing aggregated insights and benchmarks back with the private sector.

Developing New and Innovative Data Collection Methods

To build a truly predictive model, the Government should explore novel methods that provide a continuous, near-real-time view of our national cyber posture.

1. **Launch a National, Opt-In "Cyber Resilience Benchmarking" Service:**
    - The Government could sponsor a voluntary, automated service where organisations, particularly SMEs, can connect their security tools (via API) or answer a guided, interactive questionnaire. In return for sharing their anonymised performance data (e.g., alert volumes, detection times, key configuration settings), they receive a confidential report that benchmarks their resilience against their industry peers and best practices. This creates a powerful incentive for participation, provides direct value to businesses, and generates an unprecedented, live data stream for the evaluation model.
2. **Establish a "Security Telemetry Data Trust":**
    - For large businesses and critical infrastructure, the Government could spearhead the creation of a formal "data trust"—a legally and technically robust framework where multiple organisations can securely pool anonymised security operations data. This trusted environment, managed by an independent or government body, would allow for deep analysis of systemic risks, attack trends, and defensive gaps at a national scale without compromising the sensitive data of individual participants. This would be the premier source for understanding complex, cross-sector threats like supply chain attacks.

By combining the enhancement of existing surveys and reporting with the pioneering of new, voluntary, and value-driven benchmarking services, the Government can ensure its Cyber Security Policy Evaluation Model is powered by rich, relevant, and timely data. This will transform the model from a

historical record into a predictive instrument, capable of guiding policy and steering Australia towards a more secure digital future.

3. Shield-level focus for Horizon 2

3.1 Shield 1: Strong businesses and citizens

5. What could government to do better target and consolidate its cyber awareness message?

6. What programs or pilots have been successful in this context? What additional supports could be developed or scaled-up to address these issues in partnership with both education stakeholders and those with technical cyber security expertise?

Essential Supports to Develop and Scale for National Impact

Identifying successful models is the first step. Scaling them requires a deliberate, national strategy. We recommend focusing on the following areas:

1. A National Framework for Teacher Professional Development: The single greatest bottleneck to scaling these programs is teacher confidence and capability. The Government, in partnership with industry, should fund a national framework for teacher training, including micro-credentials and a repository of high-quality, co-designed teaching resources.

2. Establish a National Network of "Cyber Security PoC Hubs": To bridge the gap between classroom theory and real-world application, the Government could partner with industry to establish a network of physical or virtual "Proof of Concept (PoC) Hubs" or "Living Labs" for education. These hubs would provide a sandboxed environment where students and teachers can: ***Get hands-on experience*** *with enterprise-grade security tools (like SIEM and analytics platforms) in a safe, controlled setting.* **Simulate real-world attack and defence scenarios**, moving beyond individual puzzles to understand the full lifecycle of a security incident. * **Connect directly with industry professionals** who use these technologies daily, providing mentorship and demonstrating clear career pathways. These PoC Hubs would serve as invaluable resource centers for schools, TAFEs, and universities, making the skills needed for a career in a Security Operations Centre (SOC) accessible to the next generation.

3. Create a Public-Private "Cyber Education Fund": A dedicated, co-funded national program could provide grants directly to schools to cover participation fees for competitions, purchase specialised equipment for their local labs, and fund teacher relief days for professional development. This formalises industry's investment in building its future workforce.

4. Develop a National "Cyber Mentorship and Pathways" Platform: A centralised platform, endorsed by government and populated by industry, is needed to connect students with volunteer mentors, demystify the wide variety of careers in cyber security, and act as a clearinghouse for work experience and internship opportunities.

By systematically scaling what we know works—like CTFs—and building crucial, career-focused support structures like educational PoC Hubs, we can move from a patchwork of pilots to a cohesive national strategy that builds the resilient and talented workforce our nation needs.

7. How can Government encourage SMBs and NFPs to uptake existing cyber resources (i.e. Small Business Cyber Resilience Service, Cyber Wardens, ACNC guidance etc.)?

8. How can industry at all levels and government work together to drive the uptake of cyber security actions by SMEs and the NFP sector to enhance our national cyber resilience?

9. What existing or developing cyber security standards, could be used to assist cyber uplift for SMBs and NFP's?

10. What are the unique challenges that NFP entities face for cyber security compared to the broader business sector and what interventions from government would have the most impact in the NFP sector?

Unique Cyber Security Challenges in the NFP Sector

NFPs operate in a fundamentally different context than the broader business sector. Their challenges are not just smaller in scale; they are different in nature.

Challenge	Description
Severe Resource Scarcity	NFPs are under constant pressure to direct maximum funding to their core mission. This leaves minimal budget for dedicated cyber security personnel and enterprise-grade tools, often forcing reliance on basic, inadequate solutions.
Specialised Skill Gaps	The sector struggles to attract and retain specialised cyber security talent. IT teams, if they exist, are typically small, generalist, and overburdened. This creates a critical gap in the ability to detect, investigate, and respond to sophisticated threats.
High-Value Sensitive Data	Contrary to perception, NFPs are treasure troves of data that is highly attractive to malicious actors. This includes personally identifiable information (PII) of vulnerable clients, donor financial details, and sensitive research data.
A High-Trust, High-Risk Environment	The sector's reliance on goodwill, volunteers, and community trust makes it exceptionally vulnerable to social engineering attacks like phishing and business email compromise (BEC). Attackers exploit this trust to manipulate staff and donors.
Catastrophic Reputational Impact	A single security breach can be an extinction-level event for an NFP. The resulting loss of

Challenge	Description
	public trust can instantly cripple fundraising efforts and destroy the organisation's reputation beyond repair.

High-Impact Government Interventions for the NFP Sector

To address these unique challenges, the government should focus on interventions that build capability and provide leverage, rather than simply offering generic advice. We recommend the following:

Recommended Intervention	Rationale and Intended Impact
1. Fund "Force-Multiplier" Security Platforms	Move beyond subsidising basic antivirus and firewalls. Establish a grant program for NFPs to adopt modern, **AI-driven security analytics and automation platforms**. These solutions act as a "force multiplier," enabling small IT teams to detect complex threats, automate investigation, and manage security at a level that would otherwise be impossible. This directly addresses the resource and skills gap.
2. Promote Behaviour-Based Threat Detection	Recognise that traditional security tools, which rely on known signatures, cannot stop modern attacks that use compromised credentials. The government should champion and incentivise the adoption of **User and Entity Behaviour Analytics (UEBA)**. Such technology baselines normal activity and automatically detects deviations, spotting insider threats and credential-based attacks that are otherwise invisible. This is the most effective way to protect the high-trust NFP environment.

Recommended Intervention	Rationale and Intended Impact
3. Establish Centralised "Sector Security Hubs"	Support the creation of shared security operations centres (SOCs) or specialised Managed Security Service Providers (MSSPs) for the NFP sector. These hubs could leverage government funding to build a single, powerful security platform and offer affordable, expert monitoring and response services to hundreds of NFPs, creating an economy of scale.
4. Subsidise Modern Security Skills Training	Fund training programs that focus on modern security operations (SecOps), including threat hunting and incident response using modern analytics tools, rather than just basic cyber hygiene. This builds sustainable, in-house capability within the sector.

By focusing on these strategic interventions, the government can empower the NFP sector to defend itself effectively against the cyber threats of today and tomorrow. This approach will not only protect the vital services NFPs provide but also strengthen the security posture of the entire nation.

Exabeam stands ready to partner with the Australian Government to help realise this vision.

11. Do you consider cyber insurance products to be affordable and accessible, particularly for small entities? If not, what factors are holding back uptake of cyber insurance?

12. How well do you consider you understand the threat of ransomware, particularly for individuals and small entities? How is this threat evolving or changing?

13. How could the government further support businesses and individuals to protect themselves from ransomware attacks?

We propose the government focuses on the following high-impact interventions to shift the national mindset from static prevention to active, real-time defense.

Recommended Intervention	Rationale and Impact
1. Champion Behaviour-Based Threat Detection	The most critical gap in business defense is the failure to detect attackers using legitimate, stolen credentials to move through a network. The government should actively promote and subsidise the adoption of **User and Entity Behaviour Analytics (UEBA)**. This technology learns what is "normal" for every user and device, allowing it to automatically flag the anomalous activities—such as an administrator logging in at 3 AM from a foreign country or a user account suddenly accessing sensitive financial data—that are the tell-tale signs of a ransomware attack in progress.
2. Fund and Incentivise Security Automation	SMBs and even larger enterprises lack the human resources to manually investigate the thousands of security alerts they receive daily. The government can drive significant uplift by creating grant programs for businesses to adopt platforms that **automate threat detection, investigation, and response (TDIR)**. Automation acts as a "force multiplier," enabling small teams to instantly triage alerts, piece together the full attack timeline, and contain threats in minutes, not days. This directly neutralises the operational advantage of ransomware gangs.

Recommended Intervention	Rationale and Impact
3. Promote Centralised "Single Source of Truth" Platforms	A key failure point for businesses is having dozens of disconnected security tools that don't talk to each other. Government guidance and funding should encourage the adoption of modern, analytics-driven **Security Information and Event Management (SIEM)** platforms. These systems provide a single pane of glass, correlating events from across the entire digital estate (cloud, on-premise, endpoints) to provide a holistic view of an attack, rather than disconnected pieces of a puzzle. This is essential for understanding and stopping a complex intrusion chain.
4. Foster Actionable, Machine-Readable Threat Intelligence	The data collected from the mandatory ransomware payment reporting and the CRI is a strategic asset. The government should invest in a framework to sanitise and share this intelligence not just as reports, but as **actionable, machine-readable threat feeds**. These feeds could be directly ingested by security platforms within businesses, automatically updating their defense posture with the latest tactics, techniques, and procedures (TTPs) used by ransomware groups targeting Australia.

By implementing these measures, the government can help Australian businesses move beyond simply preparing for ransomware and towards a future where they can actively see, understand, and stop these attacks in their tracks.

14. Have you experienced or researched any vulnerabilities or impacts from cyber security

incidents that disproportionately impact your community, cohort or sector? If so, what were the vulnerabilities and impacts that your community faced?

As a global security operations leader, Exabeam has a direct view into the cyber security incidents affecting thousands of organisations across all sectors, including critical infrastructure, finance, healthcare, and government. Our research consistently shows that the most damaging attacks do not exploit exotic, zero-day software flaws. Instead, they prey on a handful of fundamental, systemic vulnerabilities that disproportionately impact organisations of all sizes.

The core vulnerability is the inability to see and understand the behaviour of users and devices *after* they have successfully logged into a network. Modern attacks are not about breaking down the front door; they are about stealing a key, walking in quietly, and pretending to be a legitimate employee.

Below are the most critical vulnerabilities we observe daily and their devastating impacts.

Core Vulnerability	**Disproportionate Impact & Real-World Scenario**
Compromised Credentials	This is the single most common attack vector. An attacker obtains valid login credentials (via phishing, malware, or purchase on the dark web) and simply logs in. **Impact:** Because the attacker appears to be a legitimate user, they bypass billions of dollars invested in perimeter defenses. They can then move freely within the network to locate and steal sensitive data, culminating in ransomware deployment. For sectors like healthcare or legal services, this leads to the catastrophic breach of highly sensitive personal records.
Insider Risk (Accidental and Malicious)	The vulnerability is a trusted user with legitimate access to sensitive systems. **Impact:** An *accidental* insider might click a phishing link, unknowingly granting an attacker access. A *malicious* insider might abuse their access to

Core Vulnerability	Disproportionate Impact & Real-World Scenario
	steal intellectual property or client data before leaving for a competitor. This disproportionately affects high-trust, high-stress environments where security controls are often seen as a hindrance to operations, leading to massive financial and data loss from a single trusted account.
Lack of Visibility and Security Tool Sprawl	Organisations have purchased dozens of disconnected security products (for email, for endpoints, for cloud). The vulnerability is that these tools do not communicate, creating critical blind spots. **Impact:** A security team may see a minor alert on a laptop, another on a server, and a third in a cloud application, but be unable to connect the dots to see they are all part of one slow, methodical attack. This allows attackers to dwell undetected for weeks or months while they map the network and exfiltrate data before the organisation even knows it has been breached.
Manual Investigation and Alert Fatigue	Security teams are overwhelmed by thousands of alerts per day from their disconnected tools, most of which are false positives. The vulnerability is the slow, manual process required to investigate each one. **Impact:** This creates a state of "alert fatigue" where critical alerts are missed. Even when a real threat is investigated, it can take hours or days to manually piece together the evidence. By then, the ransomware has been deployed or the data

Core Vulnerability	Disproportionate Impact & Real-World Scenario
	has been stolen. This speed mismatch gives a massive advantage to automated, fast-moving attackers.

The common theme across these incidents is that they are not failures of prevention, but failures of detection and response. The impacts are most severe for organisations that lack the ability to see the subtle behavioural clues of an attack in progress. Addressing these core vulnerabilities requires a strategic shift towards integrated, behaviour-based security platforms that can automate detection and response.

15. How can support services for victims of identity crime be designed to be more effective in the context of increasing demand? and

The current model of victim support for identity crime, while critical, is facing an unsustainable demand driven by a constant stream of large-scale data breaches. Designing more effective support requires a fundamental strategic pivot: we must treat the disease, not just the symptoms.

The most effective "support service" the government can offer is to prevent Australians from becoming victims in the first place. This means aggressively addressing the root cause: the failure of organisations to stop the breaches that leak personal data. By stopping the supply of stolen credentials, we reduce the demand on recovery services.

Our recommendations, therefore, are two-fold: first, focus on proactive prevention at the source, and second, leverage technology to create a scalable and empowering recovery ecosystem for existing and future victims.

A Two-Pronged Strategy for a Resilient National Identity Ecosystem

Recommended Intervention	How Technology Drives Effectiveness and Scale
1. Proactive Prevention: Stop the Breaches that Fuel Identity Crime	The most impactful way to support victims is to ensure they never become victims. The government must incentivise and mandate that organisations holding Australian data adopt **AI-driven security operations platforms**. **User and Entity Behaviour Analytics (UEBA)** technology can detect an attacker using stolen credentials inside a network *before* they can steal the data. By stopping the mass exfiltration of personal information, we directly cut off the supply chain for identity criminals. This is the only truly scalable long-term solution.
2. Early Warning System: An Automated National Breach Alert Service	Currently, victims often discover the crime months after the fact. Technology can create a national system where, upon detecting a breach, an organisation is required to report the compromised identities to a central service (an expansion of the Credential Protection Register's concept). This service would then **automatically and proactively notify affected individuals** via secure channels (e.g., myGov), providing immediate, specific guidance on what to do next. This moves the burden of discovery from the individual to an automated system.
3. Centralised Recovery: A "One-Stop-Shop" Identity Recovery Platform	Victims currently spend dozens of hours navigating a complex web of banks, telcos, and government agencies. Technology can solve this. The government should facilitate the creation of a **single, secure online platform**

Recommended Intervention	How Technology Drives Effectiveness and Scale
	for identity crime recovery. A victim could log in, report the crime once, and the platform would use APIs to automatically notify all relevant entities—placing blocks on accounts, issuing fraud alerts with credit bureaus, and flagging identity documents. This would automate the most time-consuming aspects of recovery.
4. Empowered Self-Service: AI-Powered Guidance and Support	To manage the high volume of support requests, the recovery platform should incorporate an **AI-powered virtual assistant**. This assistant could provide 24/7, personalised guidance to victims, walk them through complex steps, help them draft necessary letters, and answer common questions. This reserves human support staff for the most complex and sensitive cases, allowing the system to serve millions of Australians efficiently and effectively.

By adopting this dual approach—aggressively preventing the source of identity crime while providing powerful, automated tools for recovery—the government can build a system that is not only more effective and scalable but also one that truly empowers Australians to protect and reclaim their digital identities.

16. Which regulations do you consider most important in reducing overall cyber risk in Australia?

While all listed regulations play a role, we consider the following to be the most important models for shaping Australia's future regulatory framework, as they are inherently outcome-focused.

1. **APRA Prudential Standard CPS 234: Information Security**
   This is the gold standard for effective cyber security regulation in Australia. Its importance lies in the requirement for a regulated entity to "maintain an information security **capability** commensurate with the size and extent of threats to its information assets."
   The focus on "capability" is profound. It forces organisations to ask not "What tools do we have?" but "**What can we actually do?** Can we see an attack? Can we understand it? Can we stop it fast enough?" This directly drives investment in advanced threat detection, investigation, and response (TDIR) systems, which are essential for stopping modern attacks. It rightly treats security as a dynamic, ongoing practice rather than a static state of compliance.
2. **Company Directors' Duties (under the Corporations Act 2001)**
   This is arguably the most powerful, economy-wide driver of cyber resilience. The clarification that directors' duties of care and diligence apply to cyber security risk elevates the issue from the server room to the boardroom.
   This is critically important because it forces leadership to engage with cyber risk in a meaningful way. A board cannot delegate this duty by simply funding a tool; it must ensure the organisation has a robust and tested capability to manage the risk. This top-down pressure is essential for fostering a true culture of security and ensuring that investments are made in solutions that deliver measurable risk reduction, such as modern security analytics and automation platforms.

Conclusion: A Principle to Guide Harmonisation

While the SOCI Act and Privacy Act are vital for their specific domains, their effectiveness is amplified when viewed through the outcome-focused lens of CPS 234. For example, the Privacy Act's requirement to take "reasonable steps" to protect data should be interpreted in 2025 as having a demonstrable capability to detect and stop a data breach in progress.

Therefore, the most important regulations are those that embody the principles of **capability and accountability**. To reduce overall cyber risk in Australia, the government's harmonisation efforts should seek to infuse the outcome-focused philosophy of CPS 234 and the board-level accountability of Directors' Duties into the entire regulatory framework. This will create a more secure, productive, and competitive Australia.

17. Have regulatory/compliance requirements negatively impacted the cyber maturity of your organisation? How are you currently managing these issues?

As a security technology leader, our organisation's core mission is cyber maturity, both for ourselves and for the thousands of customers we protect. However, from our vantage point serving a diverse range of regulated Australian entities, we observe a consistent and challenging paradox: **an over-emphasis on compliance can, counterintuitively, stifle genuine cyber maturity.**

The issue is not the regulations themselves, which are well-intentioned and necessary. The negative impact arises when organisations adopt a "check-box" approach, pursuing compliance as the end goal rather than treating it as a baseline for a more robust security posture.

The "Compliance Trap" and its Impact on Cyber Maturity

We see organisations fall into a "compliance trap" that negatively impacts their cyber maturity in several distinct ways:

The Compliance Trap	Consequence for Cyber Maturity
Resource Diversion	Highly skilled security analysts spend an inordinate amount of time manually gathering logs, generating reports, and preparing evidence for audits. This is time they are *not* spending on proactive threat hunting, investigating alerts, or improving security controls. The budget is spent on proving compliance, not on improving security.
A False Sense of Security	An organisation can be 100% compliant with a prescriptive standard and still be vulnerable to a modern attack. Attackers do not care about audit reports. This false sense of security leads to complacency and a failure to invest in the advanced capabilities needed to stop a real-world adversary who is using stolen credentials to move through their network.
Innovation Lag	A rigid, compliance-first mindset can discourage the adoption of new and more

The Compliance Trap	Consequence for Cyber Maturity
	effective security technologies. If a solution isn't explicitly mentioned in a dated standard, it can be difficult to get budget for it, even if it dramatically reduces risk. This locks organisations into legacy tools and approaches while attackers constantly evolve.
Focus on Prevention, Not Detection & Response	Many compliance frameworks have historically focused on preventative controls (firewalls, passwords, etc.). This leads to underinvestment in the critical capabilities of Threat Detection, Investigation, and Response (TDIR). A mature organisation knows it *will* be targeted; its resilience is defined by how quickly it can see, understand, and stop an attack in progress.

How We Help the Market Manage These Issues

Exabeam's core strategy is to help organisations break out of this compliance trap. Our platform is designed to deliver both world-class security and streamlined compliance, proving they are not mutually exclusive goals.

We help our customers manage these issues by:

1. **Automating Compliance as a By-product of Security:** Our platform automatically collects and analyses all the necessary data for security purposes. As a result, generating evidence for compliance with regulations like the SOCI Act, CPS 234, or the Privacy Act becomes a simple, automated reporting function. This frees up security teams to focus on security.
2. **Driving Outcome-Focused Maturity:** We shift our customers' focus from "Are we compliant?" to "How quickly can we detect and respond to a threat?" By using behavioural analytics to find anomalous activity and automating the entire investigation workflow, we provide the demonstrable capability that outcome-focused regulations like CPS 234 demand.

3. **Providing Unified Visibility:** Instead of trying to manage compliance across dozens of disconnected tools, we provide a single, unified platform. This simplifies both the security operations and the compliance reporting process, reducing the burden on the organisation and allowing them to focus resources on genuine risk reduction.

In conclusion, the most significant negative impact of regulation is the mindset it can create. The solution is to foster a regulatory environment that rewards security outcomes over check-box compliance, encouraging investment in modern platforms that automate compliance tasks and deliver true cyber resilience.

• 3.2 Shield 2: Safe technology

18. What are best practice examples internationally that Australia should consider for enhancing our secure technology standards and frameworks? In particular, what approach do you consider would work best for edge devices, CER and operational technology?

International Best Practices and Their Application in Australia

Technology Area	International Best Practice Example(s)	Why It Works & Lessons for Australia
Edge Devices	**1. UK Product Security and Telecommunications Infrastructure (PSTI) Act 2. EU Cyber Resilience Act (CRA)**	**Why it Works:** These frameworks are the global gold standard for shifting responsibility to the manufacturer. They legally mandate foundational "security-by-design" principles, such as banning universal default passwords, requiring a public point of contact for vulnerability reporting, and specifying minimum software support periods. **Lessons for Australia:** Australia should

Technology Area	International Best Practice Example(s)	Why It Works & Lessons for Australia
		adopt a similar, harmonised "Secure-by-Design" law for all consumer IoT and edge devices sold in the country. This creates a powerful market-based incentive for vendors to build more secure products from the ground up, protecting millions of households and small businesses at scale without relying on end-user expertise.
Consumer Energy Resources (CER)	**A Hybrid Approach: 1. Device Level:** EU Cyber Resilience Act (CRA) principles. **2. Network Level:** Principles from the **US NERC-CIP Standards** (North American Electric Reliability Corporation Critical Infrastructure Protection).	**Why it Works:** This is a layered defense strategy. The CRA principles secure the individual devices (solar inverters, batteries). The NERC-CIP philosophy, while designed for bulk power systems, establishes the crucial concept of monitoring the grid for anomalous behaviour. **Lessons for Australia:** Australia should pursue a two-tiered approach. **Tier 1:** Mandate that all CER devices meet strong "Secure-by-Design" standards. **Tier 2:** Work with energy retailers and Distribution Network Service Providers (DNSPs) to deploy

Technology Area	International Best Practice Example(s)	Why It Works & Lessons for Australia
		behavioural monitoring across their networks. This would allow them to detect systemic threats, such as a coordinated botnet attack targeting thousands of home solar systems, which no individual homeowner could ever see.
Operational Technology (OT)	**1. The IEC 62443 Series of Standards 2. The NIST Cybersecurity Framework (CSF)**	**Why it Works:** IEC 62443 is the definitive global standard specifically designed for the unique lifecycle and operational constraints of Industrial Control Systems (ICS) and OT. It focuses on practical concepts like "zones" and "conduits" to manage risk. The NIST CSF provides the overarching risk management process (Identify, Protect, Detect, Respond, Recover) into which IEC 62443 controls can be mapped. **Lessons for Australia:** The Australian government should formally endorse the **IEC 62443 series** as the national best practice standard for OT security, particularly for entities regulated under the SOCI Act. The framework's emphasis on

Technology Area	International Best Practice Example(s)	Why It Works & Lessons for Australia
		"Detect, Respond, and Recover" is critical, as it acknowledges that prevention will eventually fail in these complex environments. Policy should incentivise investment in OT network visibility and anomaly detection platforms that can identify malicious activity without disrupting fragile legacy systems.

Conclusion: A Coherent, Multi-layered Framework

The most effective and efficient path forward is to build a national framework that learns from these international leaders. By combining strong "Secure-by-Design" legislation for manufacturers with a sophisticated "Security-by-Monitoring" approach for operators of critical services, Australia can create a resilient, multi-layered defense. This ensures that as our nation becomes more connected, it also becomes more secure.

19. How should the government work with you to support consumers and end-users to be more informed about cyber security in their products and protect themselves from cyber threats?

A Strategic Partnership Framework to Protect Consumers

Proposed Collaboration	How It Works & Protects Consumers
1. Co-develop and Distribute Actionable Threat Intelligence	**How it Works:** The government (via ASD/ACSC) and Exabeam would establish a formal channel to share anonymised, high-level threat intelligence. We see the tactics used to

Proposed Collaboration	How It Works & Protects Consumers
	breach companies in real-time; you see the national intelligence picture. **How it Protects Consumers:** By combining our views, we can identify and publicise the specific attack methods being used to steal Australian consumer data. This allows the government to issue highly specific, timely warnings to industry, enabling businesses to harden their defenses against active threats, thereby preventing the breaches that lead to consumer identity theft and fraud.
2. Elevate the National Conversation Beyond Basic Cyber Hygiene	**How it Works:** The government would feature our security strategists and threat researchers in joint webinars, official publications, and industry briefings. We would provide expert commentary not on basic password advice, but on explaining *how* modern attacks unfold and *why* organisational security maturity is a consumer protection issue. **How it Protects Consumers:** This educates the public and business leaders on the real nature of modern cyber risk. It helps consumers understand that the security of the businesses they patronise is paramount, creating public demand for better corporate security and encouraging businesses to invest in it as a competitive differentiator.
3. Jointly Champion "Security by Design" and "Security by Monitoring"	**How it Works:** We would work with government bodies to develop and promote clear standards and guidance for businesses. Exabeam can provide the real-world data to show *why* security needs to be built into

Proposed Collaboration	How It Works & Protects Consumers
	products from the start ("Security by Design") and *why* continuous monitoring is essential to catch threats that slip through preventative controls ("Security by Monitoring"). **How it Protects Consumers:** This directly leads to safer products and services on the market. When businesses are held to a higher standard and have the tools to monitor their environments, they are less likely to suffer the breaches that expose consumer data.
4. Inform Policy with Real-World, Aggregated Data	**How it Works:** Exabeam can provide the government with anonymised, aggregated data on key security metrics from the Australian market—such as average attacker dwell times, the most common types of anomalous behaviour detected, and the prevalence of insider threats. **How it Protects Consumers:** This provides an evidence base for policy-making. Instead of regulating based on theory, the government can craft smarter, more effective, and less burdensome regulations that target the real-world methods attackers are using. Better policy leads to better-protected organisations, which means safer consumers.

Conclusion: A Partnership for Systemic Change

The most impactful way for the government to work with Exabeam is to leverage our deep expertise to change the entire security ecosystem. By focusing on joint intelligence, elevating the public discourse, championing modern security principles, and informing policy with real-world data, we can collectively secure the organisations that are the frontline custodians of Australian consumer data. This is the path to protecting citizens at scale.

20. What additional guidance do you or your organisation need to manage foreign ownership, control or influence risks associated with technology vendors?

From Exabeam's perspective, both as a technology vendor ourselves and as an entity that procures technology, the key challenge in managing FOCI risks lies in the **lack of actionable transparency and the difficulty of continuous verification.** Traditional risk assessments are often point-in-time, based on self-attestation, and struggle to keep pace with dynamic ownership structures and operational realities.

To effectively manage FOCI risks, we believe the government could provide additional guidance and support in the following areas:

Area of Needed Guidance/Support	Description and Impact
1. Standardised, Verifiable Disclosure Requirements	**Need:** A universally adopted, transparent framework for technology vendors to disclose their ownership structures (including beneficial ownership), key decision-makers, and relationships with foreign governments or entities. This framework must require verifiable evidence, not just self-attestation. **Impact:** This would allow procurement teams and security professionals to make informed decisions with confidence, reducing the due diligence burden and ensuring a consistent standard across the industry. It moves beyond the current patchwork of varying questionnaires.
2. Actionable FOCI Risk Intelligence	**Need:** Government-generated, anonymised, and machine-readable intelligence feeds regarding known FOCI risks, patterns of influence, or red flags associated with certain regions or entities. This should go beyond basic blacklists and provide context that allows organisations to assess dynamic risks. **Impact:**

Area of Needed Guidance/Support	Description and Impact
	This is crucial for early warning and proactive risk management. For example, if a seemingly innocuous software update from a specific vendor suddenly originates from an unusual location, this intelligence could provide vital context for investigation.
3. Best Practices for Continuous FOCI Monitoring	**Need:** Guidance on how organisations can implement continuous monitoring programs for their technology vendors, focusing on behavioural anomalies rather than just static compliance. This includes technical indicators to look for in network traffic, software updates, and vendor support activities. **Impact:** FOCI risk is not a one-time assessment; it's an ongoing threat. Guidance on operationalising continuous monitoring, potentially leveraging security analytics platforms to detect suspicious activity linked to vendor engagement, would be invaluable. This helps identify when influence might be exerted post-procurement.
4. Clear Frameworks for "Trusted Source" Verification	**Need:** A government-backed program or certification that provides a "trusted source" designation for technology products, particularly for critical components. This would involve rigorous, independent third-party assessments of code integrity, supply chain provenance, and FOCI aspects. **Impact:** This would dramatically simplify procurement for organisations, especially SMBs, by providing a clear signal of assurance. It shifts the burden of deep technical and FOCI analysis from

Area of Needed Guidance/Support	Description and Impact
	individual organisations to a centralised, expert body.
5. Support for Secure-by-Design and Transparency Mandates	**Need:** Continued government leadership in mandating "secure-by-design" principles for all technology products entering the Australian market, alongside requirements for vendors to provide comprehensive logging and telemetry capabilities that allow customers to monitor product behaviour. **Impact:** Products that are transparent and "monitorable" inherently reduce FOCI risk. If an organisation can see exactly what a piece of software is doing, where it's connecting, and what data it's accessing, it can detect anomalous behaviour that might indicate foreign influence or compromise, regardless of the vendor's stated intentions.

By providing these types of actionable guidance and intelligence, the government can empower organisations of all sizes to navigate the complex landscape of technology vendor risks with greater confidence and effectiveness, thereby enhancing Australia's overall cyber security posture.

21. How could government better work with industry to understand data access and transfer across the economy to inform policies around secure data sharing and limit data exploitation from malicious actors?

To address the question of government-industry collaboration, we propose a shift in perspective: from a traditional focus on static data protection to a modern, dynamic understanding of **data interaction and behaviour**.

Malicious actors no longer simply breach perimeters; they exploit legitimate credentials and processes to operate undetected within networks. The key to limiting this exploitation lies in understanding the context and intent behind every data access and transfer event.

We propose a collaborative framework built on three core pillars:

1. Championing a Behaviour-Based Approach to Data Security

The government can lead by championing a move beyond compliance and perimeter-based security toward a behaviour-centric model.

- **How it works:** Instead of only asking "Is this data encrypted?", we must ask, "Is this access to the data normal?". This requires establishing a baseline of normal user and system behaviour and then using advanced analytics to detect deviations that signify a threat. This approach is essential for identifying sophisticated attacks, insider threats, and the misuse of legitimate credentials.
- **Proposed Government Action:**
    - **Establish a National Framework:** Work with industry to create a voluntary framework that defines standards for behaviour-based threat detection, especially for critical infrastructure and custodians of nationally significant data sets.
    - **Incentivize Adoption:** Promote the adoption of technologies like User and Entity Behaviour Analytics (UEBA) through government grants, R&D tax incentives, or by making it a key evaluation criterion in government procurement.

2. Creating a Real-Time, Actionable Intelligence Sharing Ecosystem

Current threat intelligence sharing often focuses on known indicators of compromise (e.g., malicious IP addresses or file hashes), which are reactive by nature. Australia needs a proactive system based on sharing threat *behaviours*.

- **How it works:** An industry partner using a modern security analytics platform can identify a new attack technique—for instance, a novel combination of PowerShell commands used to exfiltrate data. This anonymized attack *pattern* or *behaviour* is far more valuable than the specific IP address used in that single attack.
- **Proposed Government Action:**
    - **Develop a Centralized Behavioural Threat Hub:** Partner with industry leaders to create a platform where anonymized threat behaviours and adversary techniques can be shared in near real-time. This hub would provide all participants with a dynamic, early-warning system against emerging attack methodologies.
    - **Facilitate Public-Private Drills:** Sponsor and facilitate regular, cross-sector "threat hunts" where government and industry teams use these shared behavioural models to simulate and defend against attacks on critical infrastructure.

3. Leveraging AI as a Defensive Force Multiplier

The discussion paper rightly identifies AI as an accelerant for threats. The most effective countermeasure is to leverage AI for defence.

- **How it works:** The sheer volume of data access and transfer events in our economy makes manual analysis impossible. AI and machine learning are essential to automate the process of baselining normal activity, detecting anomalies, and connecting disparate security alerts into a coherent incident timeline. This dramatically reduces the time from detection to response, directly limiting the window for data exploitation.
- **Proposed Government Action:**
    - **Prioritize AI in Security Policy:** Formally recognize AI-driven security analytics as a critical capability in national cybersecurity policy.
    - **Support Sovereign AI Skills:** Invest in and partner with programs that develop local expertise in applying AI to cybersecurity challenges, ensuring Australia builds a sustainable, sovereign capability to defend itself.

22. Boosting innovation and economic prosperity is enabled when data is shared with trust and not accessed exploited by malicious actors (e.g. IP theft). How does Government and Industry work together to achieve this aim in an evolving global threat environment?

To achieve the dual aims of fostering innovation and protecting its fruits, government and industry must collaborate to build a national data security strategy based on **verifiable trust**. This means shifting from a model of prevention alone to one of deep visibility and rapid, intelligent response.

We propose a partnership focused on three key areas:

1. Establishing a National "Crown Jewels" Protection Program

The most sensitive data—the core IP that drives innovation and competitive advantage—requires a level of security commensurate with its value.

- **How it works:** This program would involve a government-industry partnership to define what constitutes "crown jewel" data and establish a gold standard for its protection. This standard would be less about building walls and more about creating total visibility. The core principle is that every access, every transfer, and every modification of this critical data is monitored,

analyzed for anomalous behaviour, and auditable. This allows organizations to confidently share data, knowing that any misuse will be immediately detected.

- **Proposed Government Action:**
  - **Co-develop a "Crown Jewels" Framework:** Create a voluntary but highly encouraged framework for identifying and securing critical IP, particularly for entities in defence, technology, and research sectors.
  - **Link R&D Funding to Security Posture:** Consider making the adoption of this framework a prerequisite for receiving significant government R&D grants or innovation funding, ensuring taxpayer-supported innovation is protected from the outset.

2. Building a "Threat-Informed Defence" for Data Sharing

Trust is eroded when organizations feel they are defending against unknown threats alone. A collaborative, threat-informed defence builds collective resilience.

- **How it works:** Instead of just sharing that a breach occurred, the focus should be on sharing the *tactics, techniques, and procedures (TTPs)* used by adversaries targeting Australian IP. For example, an industry partner could anonymously share the behavioural pattern of an attacker who used legitimate cloud service credentials to slowly exfiltrate engineering diagrams. This behavioural intelligence is far more potent for pre-emptive defence than a simple IP address.
- **Proposed Government Action:**
  - **Sponsor an Anonymized TTP Clearinghouse:** Fund and facilitate a central, non-governmental body where industry can submit and subscribe to anonymized behavioural threat intelligence related to IP theft.
  - **Integrate Intelligence into Regulation:** Evolve cybersecurity regulations to be less prescriptive about specific controls and more outcome-focused, encouraging organizations to demonstrate a capacity to detect and respond to the specific, real-world TTPs relevant to their sector.

3. Promoting Security Automation as a Business Enabler

For innovation to move at pace, security must be able to keep up. Manual security processes are a drag on both security and prosperity. The only way to manage the scale and speed of modern data sharing is through automation and AI.

- **How it works:** When security teams are equipped with platforms that automate the detection of threats and provide clear, contextualized incident timelines, they can investigate potential IP theft in minutes, not weeks. This frees up security talent to focus on proactive defence and enables

the business to adopt new collaborative technologies with confidence, knowing the security team can support them.

- **Proposed Government Action:**
  - **Launch a "Secure Innovation" Initiative:** Create a public-private initiative to showcase how AI-driven security operations (SecOps) directly enable business agility and protect IP.
  - **Invest in Sovereign Automation Skills:** Support tertiary education and vocational training programs focused on security automation, orchestration, and the application of AI in cyber defence, building a workforce that can power a secure digital economy.

23. What guidance can government provide to support the safe and responsible uptake of critical and emerging technologies?

To provide effective guidance, we must first have a clear-eyed view of the risks.

The Most Serious National Security Risk: The Automation of Trust Exploitation

The single most serious national security risk from AI is not just the creation of new weapons, but the ability to **undermine and exploit the trust** that underpins our digital infrastructure, economy, and society at an unprecedented scale and speed.

Unlike traditional threats, adversarial AI will not focus on brute force. It will focus on imitation and deception. The primary risks are:

1. **Hyper-Realistic Deception at Scale:** AI can generate phishing emails, voice clones, and disinformation campaigns that are indistinguishable from legitimate communications, leading to widespread credential theft and erosion of public trust.
2. **Automated Discovery of Novel Attacks:** AI can probe networks and software to discover new "zero-day" vulnerabilities far faster than human researchers, creating a constant state of exposure for critical infrastructure.
3. **The Compromised Insider, Amplified:** This is the most insidious threat. Once an AI gains access to a network using stolen credentials, it can operate with superhuman efficiency. It can learn the "normal" patterns of the network, identify the most valuable data, and exfiltrate it in ways that mimic legitimate user activity, making it nearly invisible to legacy, rule-based security systems. It represents the perfect spy—patient, silent, and ruthlessly effective.

Guidance for a Resilient, AI-Ready Australia

Given this threat landscape, government guidance must pivot from static, compliance-based advice to promoting a dynamic, adaptive security posture. The goal is to build resilience, not just walls.

We recommend guidance focused on three core principles:

1. **Prioritize Behavioural Insight over Perimeter Control**
    - **Guidance:** Government should advise and incentivize organizations, especially those in critical infrastructure, to adopt technologies that baseline "normal" behaviour for every user and device on their network. The primary security question must shift from "Is the firewall on?" to "Is this user's activity normal, given their role, their history, and the context of their actions?"
    - **Why it Matters:** This is the only effective way to detect an AI-powered adversary using legitimate credentials. Rules and signatures are blind to this threat; only by understanding behaviour can we spot the subtle deviations that signal a compromise.
2. **Champion a National "Assume Breach" Mentality**
    - **Guidance:** Promote a national cybersecurity posture of "assume breach." This accepts that prevention will eventually fail and places the highest priority on **speed of detection and response**. Guidance should focus on capabilities that reduce "mean time to detect" and "mean time to respond," as these are the most critical metrics in limiting damage from a breach.
    - **Why it Matters:** When an AI adversary is inside a network, every second counts. Guidance that helps organizations find, understand, and evict intruders faster is a direct investment in national resilience and business continuity.
3. **Fight AI with AI: Mandate an AI-Driven Defence**
    - **Guidance:** It is imperative that government guidance explicitly states that a human-only security team cannot defend against automated, AI-driven attacks. Organizations must be encouraged to adopt AI and machine learning as a core component of their security stack to automate threat detection, triage alerts, and investigate incidents.
    - **Why it Matters:** Human security analysts are already overwhelmed. An AI-driven defense acts as a force multiplier, handling the immense data volume and allowing human experts to focus on the most critical threats. This is not a luxury; it is the only sustainable path forward.

3.3 Shield 3: World-class threat sharing and blocking

24. What could government do to support and empower industry to take a more proactive cyber

security posture to ensure the resilience of our cyber security ecosystem? What do you think Australia's proactive cyber security posture should look like for industry?

Exabeam strongly supports the call for a more proactive national cyber posture. Creating a hostile environment for adversaries requires moving beyond a reactive, defensive stance to one of active, intelligence-driven threat engagement. For industry, this represents a fundamental shift in mindset and capability.

The Vision: What a Proactive Industry Posture Looks Like

A proactive cyber security posture for industry is not defined by higher walls, but by superior visibility and speed. It assumes that adversaries may already be inside the perimeter and operationalizes the hunt for them. This posture is built on three pillars:

1. **Threat Hunting as a Standard Business Function:** Proactive organizations do not wait for alarms. They actively and continuously hunt for threats within their own networks. This involves dedicated teams using advanced analytics to search for the subtle signs of compromise, such as unusual data access patterns, credential misuse, or lateral movement, that legacy systems miss.
2. **Defence Based on Behaviour, Not Just Signatures:** A proactive posture focuses on adversary *behaviour* rather than their tools. It recognizes that while malware can be changed, the underlying tactics of attackers—how they reconnoiter, compromise credentials, and exfiltrate data—are more consistent. By baselining normal behaviour for every user and entity, organizations can instantly spot the anomalous activity that signals an attack.
3. **Automated and Intelligent Incident Response:** Proactivity is about speed. When a threat is found, the response must be swift and precise. This is only possible through automation. A proactive organization uses AI-driven platforms to automatically investigate alerts, build incident timelines, and provide security teams with a complete, actionable picture, shrinking the time from detection to remediation from days to minutes.

Empowerment: The Government's Role in Fostering Proactivity

The government is uniquely positioned to empower and accelerate industry's shift to this proactive model. We recommend the following actions:

1. **Champion and Scale Intelligence-Led Exercises:**
    - **Action:** Expand upon the success of the CORIE program. Develop and fund a broader "National Threat Hunting Framework" that provides playbooks, tools, and training for organizations of all sizes to conduct regular, internal threat-hunting exercises based on real-world adversary TTPs (Tactics, Techniques, and Procedures).

- **Impact:** This demystifies threat hunting and transforms it from an elite activity into a standard, repeatable business process, strengthening the entire ecosystem.

2. **Incentivize the Adoption of Proactive Technologies:**
   - **Action:** Use policy levers—such as targeted tax incentives, grants, and updated procurement standards—to encourage investment in the technologies that underpin a proactive posture. Specifically, this means prioritizing platforms that provide User and Entity Behaviour Analytics (UEBA), security automation, and AI-driven security analytics.
   - **Impact:** This directly accelerates the modernization of Australia's security infrastructure, moving industry beyond legacy, signature-based tools to a more effective, behaviour-based defence.
3. **Facilitate the Sharing of Behavioural Threat Intelligence:**
   - **Action:** Establish a trusted, anonymized national clearinghouse for sharing not just indicators of compromise, but *adversary behaviours*. When a company detects a new pattern of attack (e.g., a novel way of using PowerShell for data staging), that anonymized pattern should be shared instantly with all participants.
   - **Impact:** This creates a powerful, collective defence. The entire Australian industry becomes smarter and faster with every attempted attack, making the nation a less attractive target.
4. **Cultivate a New Generation of Cyber Defenders:**
   - **Action:** Partner with universities and TAFEs to develop curricula focused on the skills needed for a proactive posture: data science, behavioural analysis, AI in security, and incident response, not just network security fundamentals.
   - **Impact:** This builds the sovereign workforce required to operate and innovate within this new security paradigm.

25. Does the government need to scope and define what Australia's proactive cyber security posture should look like for industry?

Yes, providing absolute clarity on the definition and boundaries of Active Cyber Defence (ACD) is one of the most critical steps the government can take to empower industry and secure the nation. The term ACD is currently too broad, causing uncertainty that stifles proactive security measures and creates the potential for unintended escalation.

A lack of clarity forces industry into a state of cautious paralysis, hesitant to deploy advanced defensive techniques for fear of crossing a legal or ethical line. To unlock the proactive posture Australia needs, the government must formally define the permissible scope of industry action.

We propose the government provide clarity by establishing a framework built on a fundamental distinction between two concepts: **"Active Internal Defence"** and **"Offensive Countermeasures."**

We strongly recommend that the government:

1. **Formally Adopt a Lexicon:** Publish official guidance that clearly defines "Active Internal Defence" as permissible and essential, and "Offensive Countermeasures" as the exclusive and sole domain of authorized government agencies, such as the Australian Signals Directorate.
2. **Provide a "Safe Harbour" Framework:** Create a legal and regulatory safe harbour for organizations that conduct "Active Internal Defence" in good faith and according to established best practices.

26. How could government further support industry to block threats at scale?

To help industry block threats at scale, the government's support should focus on improving the *quality* of threat intelligence, promoting the *internal* technologies that use it, and automating the *process* of defence.

We propose a framework for government support based on three key pillars:

1. Evolving Threat Sharing from "What" to "How"

Effective blocking at scale depends entirely on the quality of the intelligence driving it. Currently, many threat feeds are saturated with low-fidelity indicators (e.g., IP addresses, domains) that are reactive and quickly become outdated. The government can catalyze a paradigm shift.

- **What it looks like:** Support a move from sharing *Indicators of Compromise (IoCs)* to sharing *Indicators of Behaviour (IoBs)*. An IoB is a pattern of adversary activity—for example, "a user account that accesses the finance share drive for the first time, compresses 10GB of data, and attempts to exfiltrate it to a new cloud storage domain." This behavioural pattern is far more durable and valuable than the specific IP address used in one attack.
- **Proposed Government Action:**

- **Fund a National IoB Clearinghouse:** Enhance the existing CTIS platform or create a new initiative focused on the ingestion and dissemination of anonymized, machine-readable IoBs.
- **Update Guidance:** Update threat-sharing guidance to prioritize the contribution and consumption of these high-fidelity behavioural indicators.

2. Promoting an "Internal Immune System" for Enterprises

Horizon 1 rightly focused on blocking threats at the perimeter (ISPs). Horizon 2 must focus on blocking threats *within* the enterprise. The most dangerous threats are those that use legitimate credentials to operate inside a network. The modern form of "blocking" is not just stopping a connection, but blocking a compromised user's malicious actions.

- **What it looks like:** Organizations use advanced, AI-driven security analytics to understand normal behaviour. When an anomaly is detected (like the IoB described above), an automated response is triggered. This could be blocking the user's access to critical data, disabling their account, or isolating their machine from the network. This is enterprise-level blocking at its most effective.
- **Proposed Government Action:**
  - **Incentivize Adoption:** Create specific grants or tax incentives for the adoption of technologies that enable this internal blocking capability, such as User and Entity Behaviour Analytics (UEBA) and Security Orchestration, Automation, and Response (SOAR) platforms.
  - **Showcase the Benefits:** Actively promote the benefits of technologies mentioned in the discussion paper, like Enterprise Browser Security, as part of this broader "internal immune system" strategy.

3. Democratizing Defence Through Automated Playbooks

To achieve blocking "at scale," the process must be automated and accessible to all, not just large enterprises with dedicated security teams.

- **What it looks like:** A small business consumes threat intelligence (IoBs) from the national clearinghouse. Their security platform, using pre-configured playbooks, automatically knows how to translate that intelligence into a specific, internal blocking action without human intervention.
- **Proposed Government Action:**
  - **Sponsor the Creation of National Playbooks:** Partner with industry and ISACs to co-design a library of standardized, automated response "playbooks." These would be

practical, step-by-step guides for how to respond to specific behavioural threats (e.g., "The Ransomware Containment Playbook").

- **Promote Open Standards:** Support open standards (like the Security Orchestration and Automation Language - SOAL) to ensure these playbooks are interoperable across different security vendor platforms, preventing vendor lock-in and maximizing reach.

27. How could the use of safe browsing and deceptive warning pages be amplified?

28. What more is needed to support a thriving threat sharing ecosystem in Australia? Are there other low maturity sectors that would require ISACs, and what factors, if any, are holding back their creation?

A thriving threat-sharing ecosystem is a force multiplier for national cyber defence. While the establishment of the CTIS platform and sector-specific ISACs like the HCSN Pilot are commendable foundational steps, nurturing this ecosystem requires a deliberate focus on the quality of intelligence, the value proposition for participants, and the removal of systemic friction.

Beyond Platforms: What a Thriving Ecosystem Needs

1. **A Shift from Quantity to Quality:** The primary focus must shift from simply sharing *more* data to sharing *better* data. Threat feeds of low-quality indicators (IP addresses, file hashes) are often too noisy to be actionable. The ecosystem needs to be driven by high-fidelity, contextualized intelligence, specifically **Indicators of Behaviour (IoBs)** that describe *how* an adversary operates, not just what tools they use.
2. **A Clear Value Proposition:** For any organization, particularly a small or medium-sized business, the decision to participate comes down to a simple question: "What's in it for me?" A thriving ecosystem provides more value than it asks for. This means participants must receive timely, relevant, and *actionable* intelligence that directly reduces their risk, not just a flood of raw data that overwhelms their limited security resources.
3. **Automation and Integration:** Manual threat sharing is not scalable. The ecosystem must be built on a foundation of automation, using open standards to ensure that intelligence can be ingested directly into the security tools that organizations already use. This creates a seamless loop: detection in one organization automatically informs the defences of all.

29. How can we better align and operationalise intelligence sharing for cyber security and scams prevention?

Exabeam believes that aligning the intelligence streams for cyber security and scams prevention is one of the most impactful, high-return initiatives the government can undertake. While they target different vulnerabilities—one technical, one human—they often share the same underlying infrastructure and tactics. Operationalising this alignment will move Australia from a reactive posture to one of proactive, automated disruption.

The core challenge is that these two worlds currently operate in parallel. The ACSC is the hub for technical cyber threats, while the National Anti-Scam Centre (NASC) is the hub for scam activity. A truly operationalised system requires a bridge between them.

We propose a three-part strategy: **Unify the Intake, Specialize the Distribution, and Automate the Action.**

1. Unify the Intake: A National Cyber-Scam Intelligence Fusion Hub

The government should establish a single, unified pipeline for all threat and scam intelligence. This "Fusion Hub" would act as the central brain, ingesting data from all sources (ACSC, NASC, industry, the public) into one place.

- **How it Works:** The Hub would not replace the ACSC or NASC, but rather serve as their common intelligence engine. It would require the development of a **common intelligence format** that can codify both a technical indicator (like a malicious domain) and a scam indicator (like a mule bank account number or a specific social engineering narrative).
- **Government's Role:** Mandate and fund the creation of this central Hub and the development of the common data standard. This creates the "single source of truth" necessary for alignment.

2. Specialize the Distribution: The Right Intelligence to the Right Actor

Once the intelligence is unified, it must be directed to the actors who can take specific action. The Fusion Hub's role is to translate and route the intelligence to the appropriate entity in a machine-readable format.

- **How it Works:** The system must differentiate between intelligence types and the actions they require.

Intelligence Type	Primary Audience	Desired Action
Malicious Domain/URL	Telcos, ISPs, Enterprise Security Teams	**Block** the domain at the network or enterprise level.
Phishing Kit TTPs	ACSC, Security Vendors	**Analyse** the technique and update defensive signatures.
Fraudulent BSB/Acct No.	Financial Institutions	**Flag** or **Block** transactions to the mule account.
Scam Phone Number	Telcos	**Block** the number.
New Scam Narrative	NASC, Media, Social Media Platforms	**Alert** the public and monitor platforms for the narrative.

- **Government's Role:** Define these clear lanes of responsibility and establish the legal and regulatory framework that empowers—and, where necessary, requires—these actors to take immediate action based on the intelligence received.

3. Automate the Action: An API-Driven Response Network

The final and most critical step is to "operationalise" this system through automation. This transforms intelligence sharing from a human-driven reporting exercise into a real-time, machine-speed disruption network.

- **How it Works:** The Fusion Hub would provide a secure Application Programming Interface (API). A bank's fraud system, a telco's network filter, or a company's firewall could "subscribe" to this API. When the Hub identifies a new, verified threat, it pushes the intelligence out instantly.
- **Example Scenarios:**
  - NASC identifies a new SMS scam campaign. The malicious URL is entered into the Hub. Within seconds, an API call is made, and every participating telco has blocked the URL.
  - A bank identifies a new mule account. The BSB/Account is entered into the Hub. Within seconds, an API call is made, and every other major financial institution has flagged the account in their systems.
- **Government's Role:** Fund the development and maintenance of this national API. Champion its adoption across all critical sectors and make its use a core part of the new Scams Prevention Framework Act.

By unifying intake, specializing distribution, and automating action, we can create a system where intelligence about a threat discovered by one entity becomes an automated defence for the entire nation in near real-time. This is the future of a truly aligned and operationalised national defence ecosystem.

30. Are the roles and responsibilities of government and industry clear for cyber security in a conflict or crisis scenario? What activities, such as cyber exercises, could Government undertake to make you feel better prepared to respond in a cyber conflict or crisis?

The short answer to the first question is no. At a high level, the roles are understood—government leads on national defence and intelligence, while industry operates and defends its own networks. However, in a widespread, fast-moving crisis, these high-level principles are insufficient. The operational handoffs, legal authorities, and real-time decision-making frameworks are not yet clear enough to guarantee a swift, unified response under duress.

The CrowdStrike incident was a successful test of *cooperation* in a fault scenario. A malicious conflict scenario tests *command, control, and legal authority* in a contested environment. This is a fundamentally different and more demanding challenge. To feel better prepared, we need to move from theoretical roles to practiced, crisis-ready capabilities.

Activities to Build National Crisis-Ready Capability

The government can and should lead a series of activities designed to build the "muscle memory" required for a national response. We would feel significantly better prepared if the following were undertaken:

Activity	Description	Desired Outcome
1. Integrated National Cyber Exercises	Conduct full-scale, multi-day simulations that go beyond technical teams. These exercises must involve C-suite executives, legal counsel, and public relations teams from industry, alongside their government counterparts from	**A Battle-Tested Playbook:** The output should be a concrete, jointly-developed national crisis response playbook that identifies and resolves unforeseen friction points (e.g., legal delays, communication breakdowns,

Activity	Description	Desired Outcome
	Home Affairs, ASD, and NEMA. The scenarios must simulate **cascading failures**—for instance, an attack that simultaneously disrupts a major port, a financial institution, and a telecommunications provider.	conflicting priorities) *before* a real crisis.
2. Establish "Break-Glass" Legal Protocols	In a crisis, industry may need to take extraordinary actions (e.g., blocking traffic from an entire nation, severing connections to a key supplier) that carry legal and commercial risk. Waiting for formal government direction may be too slow.	**Pre-Authorized Emergency Powers:** The creation of a clear legislative framework that provides "safe harbour" legal protections for designated critical infrastructure entities to take pre-agreed defensive actions when specific, national-level crisis triggers are met.
3. Define and Test a Unified Command Structure	During a widespread incident, who is the single source of truth and direction? Is it the National Cyber Security Coordinator, the National Emergency Management Agency (NEMA), or ASD? Industry needs a single, clear "front door" to government for coordination and intelligence.	**A Unified Cyber Crisis Command:** A publicly documented and tested command-and-control structure that clarifies the lead agency and provides a single, operational point of contact for industry to coordinate response efforts, de-conflict actions, and receive authoritative intelligence.
4. Mandate and Fund Resilient Communications	A sophisticated adversary will likely target the very communication systems we rely on (email, corporate	**A National "Out-of-Band" Network:** The establishment and regular testing of a resilient, out-of-band

Activity	Description	Desired Outcome
	networks, mobile infrastructure). If the primary communication channels are down or untrusted, coordination becomes impossible.	communication network for critical infrastructure providers and government agencies. This would ensure that even in a degraded environment, command and control can be maintained.

In conclusion, feeling "better prepared" is the result of rigorous, realistic, and integrated preparation. By moving beyond tabletop exercises to full-scale simulations and by proactively establishing the legal and command structures needed for a crisis, the government can help ensure that on Australia's worst day, our collective response is coordinated, decisive, and pre-rehearsed.

31. How could government better incentivise businesses to adopt vulnerability disclosure policies?

To effectively incentivize the adoption of Vulnerability Disclosure Policies (VDPs), the government must lead a strategic effort to transform them from a perceived liability into a clear competitive advantage. The goal is to create an environment where a published, well-managed VDP is recognized as a hallmark of a mature, trustworthy, and resilient organization.

This requires a combination of legal protection, financial motivation, and reputational reward. We propose a four-pronged incentive strategy:

Incentive Category	Specific Government Action	Desired Outcome for Business
1. Legal & Regulatory Incentives *(Making it Safe)*	**Establish a "Safe Harbour" Framework:** Legislate clear protections for organizations that create and manage a VDP in good faith. This should shield them from punitive regulatory action for vulnerabilities	**Reduced Legal Risk & Compliance Burden:** This is the most powerful incentive. It removes the primary fear that a VDP will be used against the company. It reframes the VDP as a tool for proactive

Incentive Category	Specific Government Action	Desired Outcome for Business
	discovered and remediated through their VDP, provided they act responsibly. **Integrate VDPs into Existing Regulations:** Explicitly recognize a registered VDP as a positive control and a sign of maturity within regulatory frameworks like the SOCI Act's CIRMP and APRA's CPS 234.	compliance, which can lead to smoother audits and a better relationship with regulators.
2. Financial & Commercial Incentives *(Making it Profitable)*	**Make VDPs a Prerequisite for Government Contracts:** Phase in a requirement for suppliers bidding on significant government contracts (especially in tech and critical infrastructure) to have a registered and active VDP. **Offer a "VDP Adoption Grant":** Provide small, one-time grants, particularly for Small and Medium Businesses (SMBs), to cover the initial administrative or platform costs of establishing a formal VDP.	**Direct Commercial Advantage:** This creates a powerful market driver. Having a VDP becomes a key to unlocking a major revenue stream (government work). The grant program directly removes the initial cost barrier for smaller companies, making the decision economically sound.
3. Reputational & Educational Incentives *(Making it Prestigious)*	**Create a National VDP Register:** The ACSC could host a public, voluntary register of Australian organizations that have an approved VDP. This	**Enhanced Brand & Trust:** This allows companies to publicly signal their commitment to security. Being on the national register

Incentive Category	Specific Government Action	Desired Outcome for Business
	list would become a "gold standard" directory of trustworthy businesses. **Establish a National Cyber Stewardship Award:** Create an annual, high-profile award, presented by the Minister for Cyber Security, to celebrate organizations with exemplary VDPs and responsible disclosure programs.	becomes a marketable asset and a point of differentiation. It turns a security practice into a powerful tool for building trust with customers and partners.
4. Practical & Operational Incentives *(Making it Easy)*	**Publish Standardized VDP Templates:** The ACSC should develop and publish a set of "best practice" VDP templates, tiered for different business sizes and sectors, to lower the barrier to creation. **Facilitate a VDP "Triage-as-a-Service" Marketplace:** Support the development of a marketplace where smaller companies, who lack the internal expertise to manage vulnerability reports, can easily procure trusted Australian-based services to help them triage and manage disclosures.	**Lowered Barrier to Entry:** This removes the "we don't know how" and "we don't have the people" excuses. It makes starting a VDP a simple, straightforward process, dramatically increasing the likelihood of adoption across the entire economy.

By implementing this comprehensive strategy, the government can fundamentally shift the calculus for Australian businesses. A Vulnerability Disclosure Policy would no longer be a source of fear, but a clear, rational, and strategic investment in security, reputation, and commercial success.

32. Does Australia need a vulnerability disclosure program to provide security researchers with a mechanism for safely reporting vulnerabilities?

Yes, absolutely. The creation of a national, government-backed Vulnerability Disclosure Program (VDP) is arguably the single most important step Australia can take to unlock a vast, latent resource of security expertise and fundamentally strengthen our national cyber resilience.

Currently, ethical researchers who discover a flaw in an Australian organization's systems face a toxic choice: report it and risk their career and freedom under existing computer crime laws, or stay silent and allow the vulnerability to persist. This state of affairs is a direct threat to our national security.

We must think of a national VDP as a **"digital triple zero (000)" service for cyber vulnerabilities**. It provides a safe, reliable, and always-available channel for those with the skills to report a digital "fire" without fear of being blamed for arson.

Core Components of a National VDP

A successful national program, likely administered by the Australian Cyber Security Centre (ACSC), must be built on the following pillars to be effective:

Component	Description: How It Works	Primary Benefit
1. Unambiguous Legal Safe Harbour	This is the bedrock of the entire program. The government must legislate clear and robust "safe harbour" provisions that protect researchers from civil and criminal liability when they discover and report vulnerabilities in good faith, adhering to the program's rules.	**Removes Fear & Unlocks Talent:** This single action would signal a seismic shift, transforming Australia from a legally risky environment into a globally attractive one for security research. It unlocks a vast pool of ethical talent, both domestic and international.
2. ACSC as the Triage Hub & Intermediary	The ACSC would act as the central, trusted intake point. Researchers could report any	**Provides a "Last Resort" Reporting Channel:** This solves the problem of

Component	Description: How It Works	Primary Benefit
	vulnerability to the ACSC, especially for companies that lack their own VDP. The ACSC would then confidentially contact the affected organization and facilitate a coordinated disclosure process.	discovering vulnerabilities in immature or unresponsive organizations. It ensures that no critical vulnerability falls through the cracks, and it protects the researcher from having to engage directly with a potentially hostile company.
3. Guaranteed Anonymity and Confidentiality	The program must offer a mechanism for researchers to report vulnerabilities anonymously or pseudonymously. This builds trust and encourages reports from individuals who may not want public recognition or association.	**Increases the Volume of Reports:** Many researchers are motivated by the technical challenge, not public credit. Guaranteeing their privacy encourages a higher volume of disclosures and protects them from potential retaliation.
4. Clear and Simple "Rules of the Road"	The program must publish a clear, simple set of rules for engagement. This would define "good faith" research, outlining what is in scope (e.g., identifying a flaw) and what is out of scope (e.g., exfiltrating sensitive data, performing denial-of-service attacks).	**Provides Certainty for All Parties:** Clear rules prevent accidental overstepping by researchers and provide organizations with the assurance that the program has guardrails. It minimizes the risk of a good-faith effort causing unintended harm.

By establishing this program, the government would do more than just fix a policy gap. It would initiate a profound cultural shift—moving from an adversarial relationship with the security research community to a collaborative partnership. It would make Australia a safer place to do business, a more attractive place for top security talent, and a much harder target for our adversaries. This is a foundational investment in a secure digital future.

3.4 Shield 4: Protected critical infrastructure

33. How effective do you consider the SOCI Act at protecting Australia's critical infrastructure? Are the current obligations proportionate, well-understood, and enforceable?

The *Security of Critical Infrastructure Act 2018 (SOCI Act)* is an essential and world-leading legislative foundation. It has been highly effective in establishing a national baseline for security and creating a common language for government and industry to discuss risk. However, its ultimate effectiveness will be determined by its evolution from a framework for compliance into a catalyst for genuine, demonstrable cyber resilience.

The Act's strength is in defining *what* needs to be protected and the high-level obligations for doing so. The challenge, and the focus for Horizon 2, lies in maturing the *how*—ensuring that the spirit of the law is met, not just the letter.

Here is our assessment of the current obligations:

Obligation	Proportionate?	Well-Understood?	Enforceable?
Critical Infrastructure Risk Management Program (CIRMP - Part 2B)	**Yes.** The risk-based approach is the very definition of proportionate, allowing entities to tailor their response to their specific threat environment.	**The "What," Yes; The "How," No.** The requirement to *have* a program is clear. However, how to effectively implement, monitor, and provide continuous evidence of that program's effectiveness is a significant challenge for many organizations.	**Difficult, at present.** As noted in the discussion paper, enforceability is weak without a mandatory audit scheme. True enforcement requires moving beyond paper-based assessments to evidence-based validation of security outcomes.
Mandatory Cyber Incident Reporting	**Yes.** The 12 and 72-hour reporting timelines are proportionate to the need for rapid national-	**Yes.** The obligation itself is one of the most	**Yes.** Non-compliance with reporting obligations is relatively straightforward to

Obligation	Proportionate?	Well-Understood?	Enforceable?
	level awareness during a significant incident.	clearly understood parts of the Act.	identify and enforce after the fact.
Government Assistance Powers (Part 3A)	**Yes, in principle.** In a true national crisis, these powers are a proportionate last resort. However, their use would signify a failure of earlier, more routine defensive measures.	**The Powers, Yes; The Triggers, Less So.** The powers themselves are understood, but the precise thresholds that would trigger government intervention are less clear to industry, which can create uncertainty.	**Yes.** These powers are directly enforceable by government action.

Recommendations for Maturing the Framework

To enhance the Act's effectiveness, we support the direction outlined in the discussion paper and offer the following recommendations:

1. **Champion "Evidence-Based Assurance" over "Checklist Compliance"** The single greatest step to increase the Act's effectiveness is to shift the focus from attestation to evidence. The proposed independent audit requirement is the right mechanism to achieve this. These audits should not just review paperwork; they must be empowered to seek tangible proof that security controls are implemented, tested, and effective. This requires organizations to have platforms that can provide this evidence—showing, for example, that anomalous user behaviour is being detected and acted upon.
2. **Clarify the "How" with Actionable Guidance** While the Act rightly avoids being overly prescriptive, there is a clear need for sector-specific guidance on *how* to meet the CIRMP obligations. This guidance should highlight modern security approaches, such as the importance of behavioural analytics for detecting insider threats and compromised credentials, which are key risks for critical infrastructure.
3. **Invest in Automated Compliance and Reporting** A key reason for the gap between "well-understood" and "enforceable" is the high manual effort currently required for monitoring and reporting. The government should encourage the adoption of security platforms that can

automate the collection of audit evidence and map it directly to the CIRMP requirements. This lowers the compliance burden, improves accuracy, and makes the framework more readily enforceable.

In conclusion, the SOCI Act is a strong and proportionate framework. Its current weakness lies in the gap between understanding the obligations and an organization's ability to prove its continuous adherence to them. By introducing robust, evidence-based audits and supporting industry with guidance on modern security practices, the government can close this gap and transform the Act from a baseline requirement into a powerful engine for national cyber resilience.

34. Are there significant cyber security risks that are not adequately addressed under the current framework?

While the SOCI Act provides a strong foundation for asset-based protection, several significant, systemic cyber security risks exist at the seams of the current framework or are emerging so rapidly that they warrant specific, dedicated attention. These are risks that can impact critical infrastructure indirectly or operate in ways that traditional regulatory models struggle to address.

We have identified four key areas that require a dedicated focus:

Risk Area	Description: Why It's a Gap
1. Systemic Supply Chain Concentration Risk	The CrowdStrike incident was a stark reminder that the security of Australia's critical infrastructure is critically dependent on a small number of key technology providers (e.g., cloud providers, managed security services, widely used enterprise software). The SOCI Act rightly places the onus on the asset owner to manage their supply chain risks. However, it does not directly address the **concentration risk** itself. A single successful attack on a "systemically important technology provider" could represent a cascading, one-to-many failure across numerous critical infrastructure sectors simultaneously.

Risk Area	Description: Why It's a Gap
2. Coordinated, Cyber-Enabled Trust Erosion	The SOCI Act is designed to protect the *technical operation* of an asset. It is not explicitly designed to counter a coordinated, cyber-enabled disinformation campaign intended to destroy public *trust* in that asset. An adversary could use AI-driven deepfakes, bot farms, and social media manipulation to convince the public that the financial system is collapsing or the water supply is tainted—even with no technical breach—causing panic, bank runs, or civil unrest. This is a direct attack on national resilience, using cyber tools to achieve a psychological and societal effect.
3. The Operational Technology (OT) "Security Debt"	While OT is covered under the "all hazards" approach, the framework may not fully account for the immense "security debt" inherent in these environments. Many systems that run our physical world are decades old, cannot be patched, use insecure-by-design protocols, and were never intended to be connected to IT networks. This isn't just a matter of applying a security framework; it's a systemic problem of legacy fragility that generic risk management programs may struggle to adequately mitigate without massive capital investment in modernization.
4. Autonomous, AI-Driven Adversaries	The current framework—and indeed, most global cyber defence doctrine—is predicated on responding to human-driven or tool-assisted attacks. It may be unprepared for the advent of truly **autonomous AI adversaries**. Such an

Risk Area	Description: Why It's a Gap
	agent could be given a high-level objective (e.g., "destabilize the Victorian energy grid") and independently discover vulnerabilities, chain exploits together, and adapt its tactics at machine speed to evade detection. This class of threat could operate faster than any human-in-the-loop defence and response cycle, rendering current incident response plans obsolete.

Moving Forward

Addressing these risks does not necessarily require rewriting the SOCI Act. Rather, it requires a broader view of national resilience that builds upon the Act's foundation:

- **For Supply Chain Risk:** Begin identifying and engaging with "Systemically Important Technology Providers" to co-design a framework for transparency and evidenced-based assurance.
- **For Trust Erosion:** Deepen collaboration between the ACSC, NEMA, and agencies responsible for countering foreign interference and disinformation to develop a national response playbook for this specific threat.
- **For OT Security Debt:** Consider targeted national initiatives or funding programs to accelerate the modernization of the most critical and fragile OT systems.
- **For AI Adversaries:** Champion the development and adoption of AI-driven *defensive* systems as the only realistic countermeasure, and heavily invest in R&D and exercises that simulate and defend against this emerging threat class.

The SOCI Act is an excellent shield for the threats we know. The next challenge is to build the capabilities needed to defend against the threats that will define the coming decade.

35. Is the regulatory burden on industry proportionate to the risk and outcomes being sought?

In its current form, the regulatory burden of the SOCI Act *can be* disproportionate **if organizations attempt to meet it with outdated, manual processes**. However, the burden *is* proportionate **when organizations meet it with modern, automated security platforms**.

The issue is not the existence of the obligations, but the *method* of compliance.

The Two Faces of Regulatory Burden

We see a clear dichotomy in how the burden is experienced. It depends entirely on an organization's approach to security and compliance.

Aspect	**Manual, Legacy Approach** (Disproportionate Burden)	**Automated, Modern Approach** (Proportionate Burden)
Effort & Cost	Extremely high. Requires teams of staff manually collecting logs, writing reports, and preparing for audits. Costs are operational expenses that do not improve security.	Low to medium. The platform automates data collection, correlation, and reporting. Costs are an investment in a capital asset that directly improves security.
Nature of Work	Reactive and periodic. A massive, stressful effort is required before an audit or when an incident occurs.	Proactive and continuous. The system provides "always-on" assurance and is audit-ready at all times.
Outcome	"Compliance Theatre." Produces a point-in-time snapshot that proves paperwork is in order but may not reflect genuine security posture.	**Demonstrable Resilience.** Provides real-time evidence of effective security controls and the ability to detect and respond to threats.
Risk Mitigation	Low. The manual effort is so high that it detracts from actual	High. The activities required for compliance are the *same*

Aspect	Manual, Legacy Approach (Disproportionate Burden)	Automated, Modern Approach (Proportionate Burden)
	threat hunting and security improvement activities.	*activities* that detect and stop cyber attacks.

When compliance is a manual, backward-looking exercise, the burden is immense and the security outcome is minimal. This is disproportionate. When compliance is the automated output of a proactive, forward-looking security program, the burden is simply the cost of doing secure business, which is entirely proportionate to the risk.

36. What support would assist critical infrastructure owners and operators to mature their cyber and operational resilience practices? What role should government play in enabling uplift, including through tools, guidance or incentives?

The maturation of cyber and operational resilience is the central challenge of Horizon 2. For critical infrastructure owners and operators, this journey moves beyond compliance and into the complex realm of building a deeply ingrained security culture and capability. To be successful, the government's role must evolve from that of a regulator to that of a **national resilience partner**.

This partnership should be built on a clear strategy of reducing friction and increasing capability. The government can provide the most effective support by delivering a coordinated program of practical tools, actionable guidance, and powerful incentives.

Government Role/Lever	Specific Government Action	Desired Outcome for Industry
Guidance: From 'What' to 'How'	**1. Develop National Resilience Playbooks:** The ACSC should partner with industry to co-author and maintain a library of practical, step-by-step playbooks for responding to specific, high-risk scenarios (e.g., "The	**Clarity and Direction:** This answers the question, "What does 'good' look like?" It gives CISOs and boards a tangible roadmap for investment and effort, moving maturity from an

Government Role/Lever	Specific Government Action	Desired Outcome for Industry
	Ransomware Containment Playbook," "The Compromised Insider Playbook"). **2. Publish a National Cyber Maturity Model:** Create a clear, five-level maturity model for critical infrastructure. This would allow an organization to benchmark its current state and provide a clear, aspirational roadmap for improvement.	abstract concept to a measurable goal.
Tools: Democratizing Advanced Capability	**1. Establish a National Threat Hunting Platform:** Provide a sandboxed, virtual environment where organizations can test their defences against simulated attacks based on real-world threat intelligence. **2. Provide an "IoB" Threat Feed:** Enhance the existing CTIS platform to provide a high-fidelity feed of actionable **Indicators of Behaviour (IoBs)**, not just basic Indicators of Compromise. This tells defenders *how* adversaries are operating.	**Access and Empowerment:** This gives all critical infrastructure operators—regardless of size—access to advanced tools that are typically only available to the largest banks or intelligence agencies. It democratizes the ability to practice and prepare.
Incentives: Aligning Security with Business Goals	**1. Create a "Resilience Modernization Fund":** Establish a targeted tax	**Reduced Friction and Clear ROI:** Incentives directly address the primary barriers to

Government Role/Lever	Specific Government Action	Desired Outcome for Industry
	incentive or grant program specifically for the adoption of modern security technologies that automate compliance and improve threat detection (e.g., UEBA, SOAR platforms). **2. Offer "Safe Harbour" for Information Sharing:** Introduce clear legislation that protects organizations from liability when they share threat intelligence in good faith through registered ISACs. **3. Streamline Procurement for Assured Solutions:** Create a program to certify security solutions that meet the highest standards for automation and evidence-based assurance, giving them a preferred status in government and critical infrastructure procurement.	investment: cost and fear. They make the business case for maturing security posture undeniable by lowering the financial risk and removing legal uncertainty, turning security from a cost centre into a strategic investment.

In conclusion, the most effective support the government can provide is a holistic program that makes the path to maturity the path of least resistance. By providing clear roadmaps, powerful shared tools, and compelling financial and legal incentives, the government can foster an environment where maturing from a compliant entity to a resilient one is not just a regulatory obligation, but a clear, achievable, and logical business strategy for every critical infrastructure operator in Australia.

37. How can the Australian Government support private sector partners to better engage with government security requirements, including certifications and technical controls?

The most effective support the government can provide is to champion a modern framework for engagement that values **dynamic evidence over static documents** and **demonstrable outcomes over checklists**.

Here are four key suggestions on how the government can facilitate this evolution:

1. Suggestion: Shift from Static Reports to "Live" Evidence

Current Engagement: The current model often relies on partners producing lengthy, static PDF documents for security assessments. This is incredibly time-consuming, costly, and provides a picture of security that is outdated the moment it's printed.

A Better Engagement Model: The government can spearhead a shift by encouraging partners to demonstrate compliance through live dashboards and real-time queries.

- **How it Works:** Instead of a static report, an assessor would be given secure, read-only access to a partner's centralized security analytics platform. To verify a control, they could ask, "Show me all attempts to access disabled accounts in the last 30 days." A modern platform can provide this definitive answer in seconds, pulling data from hundreds of sources across cloud and on-premises environments.
- **Government's Role:** Champion this "live audit" approach. Develop standards for secure, role-based access for assessors and recognize it as a high-maturity method for demonstrating compliance.

2. Suggestion: Prioritize "Detection-in-Depth" as a Core Capability

Current Engagement: Assessment of technical controls often focuses on validating that a preventative control is *present* (e.g., "Is MFA enabled?"). It less frequently asks the more critical question: "How do you detect when this control is bypassed or fails?"

A Better Engagement Model: Engagement and assessment processes should evolve to weigh detection and response capabilities as heavily as prevention.

- **How it Works:** A partner would demonstrate not only that they restrict administrative privileges, but that they can instantly detect when a legitimate admin account is compromised. This requires platforms with built-in **User and Entity Behaviour Analytics (UEBA)**. This capability automatically baselines normal activity for every user and system, then flags high-risk deviations, providing an essential security layer that proves controls are effective.

- **Government's Role:** Update assessment criteria within frameworks like IRAP to include mandatory testing of a partner's ability to detect common attack behaviours that bypass preventative controls.

3. Suggestion: Provide Clear Implementation Guidance for Modern Architectures

Current Engagement: Applying frameworks like the ISM to modern, multi-cloud SaaS environments can be ambiguous. Partners spend significant time and resources interpreting how legacy controls apply to ephemeral containers or serverless functions.

A Better Engagement Model: The government can publish official "implementation profiles" or "overlays" for the ISM that are tailored to specific modern technologies (e.g., AWS, Azure, Google Cloud).

- **How it Works:** This guidance would clarify government expectations on how logging and monitoring controls should be implemented in a distributed cloud environment. This enables partners to confidently configure their security monitoring platforms to collect the right data and apply relevant detection rules for that specific architecture.
- **Government's Role:** Work with cloud providers and security vendors to develop and maintain these implementation profiles, ensuring they remain current as technology evolves.

4. Suggestion: Streamline Assurance through Automated, Continuous Reporting

Current Engagement: The "certify once, re-certify later" model is inefficient and provides no assurance about a partner's security posture in the long intervals between audits.

A Better Engagement Model: The government could pilot a "Trusted Partner Network" based on the principle of continuous assurance.

- **How it Works:** Vetted partners could securely stream anonymized security posture and compliance data directly from their security platforms to a government portal. This leverages the capabilities of modern security tools built for automated evidence collection and API-based integration. A dashboard could show that a partner has maintained 99.9% compliance with key controls over the last quarter.
- **Government's Role:** Define the standards and secure API for this data exchange. This would create a high-trust environment, reduce audit fatigue for partners, and provide the government with unprecedented, near real-time visibility into the security of its supply chain.

By adopting these engagement models, the government can lower the barrier for innovative partners, increase the quality and reliability of security assurance, and build a more dynamic and truly collaborative security partnership.

38. How are Australian Government security requirements or frameworks being considered or adopted among private sector partners, including in critical infrastructure?

• 3.5 Shield 5: Sovereign capabilities

39. What role should government play in supporting the development and growth of Australia's cyber workforce? What initiatives, pilots or policy ideas do you think would best support industry to grow?

The government's most effective role in growing Australia's cyber workforce is to act as a **strategic enabler**. While industry must lead in hiring and training, the government is uniquely positioned to address the systemic barriers and market failures that currently constrain the talent pipeline. The goal should be to make the cyber profession more **visible, accessible, and attainable** for all Australians.

We propose a program of initiatives focused on three key areas: **Redefining the Pathways, Simplifying the Standards, and Incentivizing the Transition.**

Strategic Area	Proposed Initiative/Pilot	How it Supports Industry Growth
1. Redefine the Pathways *(Beyond the University Degree)*	**Launch a "National Cyber Apprenticeship Program":** A government-co-funded program where industry partners take on apprentices for paid, on-the-job training in specific cyber roles (e.g., Security Analyst, Threat Hunter). The government's	**Unlocks New Talent Pools:** These initiatives directly address the industry's struggle to find talent outside of the traditional, highly competitive university graduate market. They create a structured, low-risk pathway for companies to hire and mould talent with

Strategic Area	Proposed Initiative/Pilot	How it Supports Industry Growth
	contribution would offset the initial training cost and salary burden for the employer. **Pilot a "Return-to-Cyber" Fellowship:** A program specifically for mid-career professionals, particularly women returning to the workforce. It would combine a short, intensive technical bootcamp with a guaranteed internship at a partner company.	diverse life experiences and proven professional skills.
2. Simplify the Standards *(Navigating the "Certification Maze")*	**Establish a "National Cyber Skills Framework":** The government, through the ACSC, should lead the creation of a definitive framework that maps common cyber job roles to specific, verifiable skills and competencies—independent of any single vendor certification. **Create a "Certification Equivalency Register":** This register would formally recognize how different industry certifications (e.g., CompTIA, ISC², GIAC) map to the national skills framework. This would allow an employer to instantly understand the	**Reduces Hiring Friction and Cost:** This tackles a major pain point for industry. It allows hiring managers to focus on what a candidate can *do*, not just what certifications they have. It broadens the candidate pool and makes it easier to assess the true skills of applicants from non-traditional backgrounds.

Strategic Area	Proposed Initiative/Pilot	How it Supports Industry Growth
	value and relevance of a candidate's qualifications.	
3. Incentivize the Transition *(Making Cyber an Attractive Career Change)*	**Offer a "Cyber Skills Transition Rebate":** Provide a direct-to-citizen financial rebate (e.g., up to $5,000) for individuals who complete an accredited cyber security training course and successfully transition into a cyber role. This helps offset the cost of re-skilling. **Fund a National "Day in the Life" Campaign:** A high-profile public awareness campaign, led by the ACSC and industry partners, that showcases the true diversity of roles in cyber security. It should feature real people—from threat hunters and data scientists to GRC analysts and ethical hackers—to break down the "hacker in a hoodie" stereotype.	**Drives Market Demand and Widens the Funnel:** The rebate makes the financial decision to re-skill much easier for individuals, directly increasing the supply of new talent for industry to hire. The awareness campaign helps solve industry's perception problem, attracting people from diverse backgrounds (e.g., problem-solvers from law, analysts from finance, investigators from law enforcement) who may never have previously considered a career in cyber.

By focusing on these strategic initiatives, the government can do more than just support industry—it can actively shape and grow the market for cyber talent. This partnership will ensure Australia has the skilled, diverse, and sustainable workforce required to secure our digital future.

40. What have been the most successful initiatives and programs that support mid-career

transitions into the cyber workforce and greater diversity in technology or STEM-fields more broadly?

41. What are some of the industries with highly transferrable skill sets that could be leveraged to surge into the cyber workforce? Is there any existing research/data that could support these efforts?

To rapidly "surge" into the cyber workforce, Australia must look beyond traditional IT pathways and systematically target industries rich in professionals who already possess the core *aptitudes* and *mindsets* essential for cyber security. Technical skills can be taught relatively quickly; analytical rigor, an investigative mindset, and the ability to perform under pressure are far more ingrained.

The most successful transitions will come from industries where professionals are already trained to manage risk, identify anomalies, and follow complex processes.

Industries with Highly Transferrable Skills

The following table outlines key industries and maps their inherent skills to specific, in-demand cyber security roles:

Source Industry	Key Transferable Skills & Attributes	Target Cyber Security Roles
Finance & Banking	**Fraud Analysis, Anti-Money Laundering (AML), Risk Assessment:** Deep experience in pattern recognition, data analysis, regulatory compliance, and identifying deceptive financial behaviour.	Security Operations Center (SOC) Analyst, Threat Intelligence Analyst, Governance, Risk & Compliance (GRC) Analyst, Cyber Fraud Analyst.
Law Enforcement, Intelligence & Military	**Investigative Process, Evidence Handling, Adversarial Thinking:** A	Digital Forensics and Incident Response (DFIR) Analyst, Threat Hunter, SOC Analyst

Source Industry	Key Transferable Skills & Attributes	Target Cyber Security Roles
	natural aptitude for investigation, analytical discipline, working with incomplete information, and understanding an opponent's motives.	(Tier 2/3), Physical Security Convergence Specialist.
Legal & Compliance	**Policy Interpretation, Logical Reasoning, Risk Management:** Expertise in deconstructing complex regulations, building logical arguments, assessing risk, and communicating precisely.	GRC Analyst, Cyber Policy Advisor, Privacy Officer, Security Auditor.
IT & Network Operations	**System Architecture, Network Protocols, Troubleshooting:** Foundational knowledge of how systems work, what "normal" traffic looks like, and how to diagnose complex technical problems.	Security Engineer, Cloud Security Specialist, Network Security Analyst, Security Architect.
Customer Service & Technical Support	**Communication Under Pressure, Pattern Recognition, De-escalation:** Front-line experience in methodical problem-solving, identifying recurring issues from user reports, and guiding	SOC Analyst (Tier 1), Phishing Response Analyst, Security Awareness & Training Specialist, Identity and Access Management (IAM) Administrator.

Source Industry	Key Transferable Skills & Attributes	Target Cyber Security Roles
	people through stressful situations.	
Project & Program Management	**Stakeholder Management, Risk Mitigation, Process Optimization:** Proven ability to manage complex projects, communicate with diverse stakeholders (from technical to executive), and manage budgets and timelines.	Cyber Security Project/Program Manager, Security Engagement Lead, GRC Program Manager.

Supporting Research and Data

While there may not be a single, definitive report on this topic, a strong body of evidence can be assembled from several authoritative sources to support these efforts. Government should leverage and promote this existing work:

1. **Professional Cyber Security Bodies:**
    - **Source:** Organizations like **(ISC)², ISACA, and CompTIA** regularly publish global workforce studies.
    - **Data:** These reports consistently highlight the most in-demand skills and the widening gap between supply and demand. They provide data that can be used to map the skills outlined above to the specific roles that industries are desperate to fill. For example, their data validates the huge demand for analytical and communication skills, not just technical prowess.
2. **Major Consulting and Technology Firms:**
    - **Source:** Firms such as **Deloitte, PwC, McKinsey, and major technology companies** frequently publish research on the "future of work," digital skills, and workforce transformation.
    - **Data:** These reports often contain analyses of skills adjacencies—identifying which non-tech roles have skills that are highly compatible with technology fields. This data can provide a quantitative foundation for targeting specific professions for reskilling initiatives.
3. **Government Skills Frameworks (International):**

- **Source:** International government bodies, such as the **U.S. National Institute of Standards and Technology (NIST)** with its "NICE Cybersecurity Workforce Framework," and the **UK's Cyber Security Skills Framework.**
- **Data:** These frameworks break down cyber security into specific tasks, knowledge areas, and skills. They provide a powerful tool for "deconstructing" a cyber role into its component parts, which can then be matched against the skills prevalent in the target industries listed above. This provides a methodical, data-driven approach to identifying transition pathways.

4. **Veteran Transition Studies:**
   - **Source:** Numerous studies and reports from organizations focused on veteran employment.
   - **Data:** These provide powerful qualitative and quantitative evidence of the successful transition of military personnel into cyber security roles, validating the transferability of skills like discipline, intelligence analysis, and operational planning.

The government's role should be to aggregate this existing research, use it to validate the targeting of these industries, and partner with industry associations to build dedicated outreach and transition programs for these high-potential talent pools.

42. How can industry, academia, think tanks and government best work together to set research priorities and drive innovation to further our strategic, economic and community interests and achieve our common goals?

43. How can government and academia enhance its partnership and promote stronger people-to-people links and collaboration on research and policy development activities?

44. How would we best identify and prioritise sovereign capabilities for growth and development across government and industry?

45. What are the areas of most concern for ICT concentration and what do you consider would be most effective as mitigation strategies to explore?

Yes, we view attributions, advisories, and sanctions as **effective and essential foundational tools**, but they are not a complete strategy in themselves. They are the necessary first steps in building a credible deterrence posture. Their primary effectiveness lies in creating a baseline of accountability and providing defenders with critical, actionable intelligence.

Here is our assessment of the current toolkit:

Tool	Primary Effectiveness	Key Limitation
Attributions	**Norm-Building & Shaming:** Publicly naming state actors helps build an international consensus that such activity is unacceptable. It imposes a reputational cost and rejects any notion of plausible deniability.	Can be slow to assemble evidence for and may have little impact on shameless actors who are immune to reputational damage.
Advisories	**Collective Defence:** Sharing the TTPs of malicious actors is a highly effective defensive tool. It empowers the entire ecosystem to protect itself, raising the cost and effort for the attacker on their next operation.	This is a purely defensive measure; it does not directly impose a punitive cost on the adversary or deter their future intent.
Sanctions	**Direct Economic Cost:** When targeted effectively, sanctions can disrupt the financial operations of cybercriminal groups and the individuals who	Can be symbolic if the target holds no assets within the reach of the sanctioning nations. State actors can often find ways to circumvent them.

Tool	Primary Effectiveness	Key Limitation
	support them, creating tangible friction and imposing direct costs.	

While these tools are crucial, to effectively combat the rising threats in Horizon 2, Australia must evolve its strategy from one of *reactive cost imposition* to one of *proactive deterrence by denial and entanglement*. We propose Australia considers developing the following tools of cyber diplomacy and deterrence, leveraging advanced security capabilities as key enablers:

Proposed Tools for Horizon 2, Enabled by Advanced Cyber Capabilities

Strategic Category	Proposed Tool/Initiative	How a Modern Security Analytics Platform Enhances This
1. Proactive Diplomatic & Norm-Building	**Champion a "Safe Harbour for Critical Infrastructure" Norm:** Australia should lead a diplomatic initiative to establish a clear international norm that targeting the critical infrastructure of another nation (e.g., healthcare, energy grids) during peacetime is an unacceptable violation of state sovereignty, akin to a physical attack. **"Exporting Defence" Initiative:** A formal program where the Australian Cyber Security Centre (ACSC) provides structured training, capacity building, and intelligence sharing to partner nations, particularly in	A platform that monitors *all* activity within critical infrastructure environments, especially user and entity behavior, can provide the continuous evidence needed to demonstrate adherence to strict security norms and rapidly detect any deviations or compromises. This provides the auditable proof of diligence required for such a norm. A robust security analytics platform, which provides comprehensive visibility into network activity and sophisticated anomaly detection, is fundamental to building strong national cyber

Strategic Category	Proposed Tool/Initiative	How a Modern Security Analytics Platform Enhances This
	Southeast Asia and the Pacific, to help them build their own national cyber defence capabilities.	defense capabilities. By providing access to such platforms, Australia can enable partner nations to effectively monitor their own networks, detect advanced threats, and respond proactively, thereby lifting the collective regional security posture.
2. Economic & Industrial Deterrence	**Form a "Trusted Technology Alliance":** A multilateral alliance of like-minded nations that agree to a common, high standard for security and transparency in their technology supply chains. Membership would confer a significant trade and investment advantage. **Coordinated Vulnerability Disclosure:** When a government or private sector entity in an allied nation discovers a significant vulnerability in software or hardware originating from an adversarial state, the alliance would coordinate a responsible, public disclosure.	This requires internal vigilance over how allied technologies are being used and if they are exhibiting any anomalous behavior. A platform that can ingest and analyze logs from various technology layers (network, endpoint, cloud, applications) helps continuously verify the integrity and secure operation of these trusted technologies within an organization's environment. It allows for the detection of any misuse or unexpected activity from "trusted" components. While the platform itself doesn't find vulnerabilities, it's crucial for defense once one is disclosed. A security analytics platform enables rapid

Strategic Category	Proposed Tool/Initiative	How a Modern Security Analytics Platform Enhances This
		detection of *exploitation* attempts once a zero-day or N-day vulnerability in adversarial technology is announced. Its ability to correlate disparate events and identify unusual patterns can quickly flag if an organization is being targeted or compromised via that specific vulnerability, allowing for swift mitigation.
3. Active, Collective Deterrence	**Adopt a "Persistent Engagement" Posture (with Allies):** Publicly declare a policy, in concert with Five Eyes or AUKUS partners, of "persistent engagement." This means proactively hunting for adversary activity on our own networks to detect and disrupt campaigns *before* they achieve their objectives. **Conduct Public, Joint "Threat Hunt" Exercises:** Australia and its allies should conduct regular, joint cyber exercises where the "Red Team" explicitly simulates the TTPs of a named, real-world state actor (e.g., APT40, as mentioned in the context).	This is a direct alignment. A security analytics platform with advanced threat hunting and **behavioral analytics capabilities** is the core engine for persistent engagement. It allows defenders to proactively search for subtle indicators of compromise by looking for deviations from normal user and system behavior, rather than waiting for an alert. This capability is essential for uncovering stealthy, sophisticated attacks that bypass traditional defenses. These exercises rely on the ability to collect, analyze, and correlate vast amounts of

Strategic Category	Proposed Tool/Initiative	How a Modern Security Analytics Platform Enhances This
	The results and defensive lessons should be publicly shared.	security data to identify simulated adversary activity. A modern security analytics platform provides the necessary tools for participants to perform these hunts effectively, visualize attack chains, and test their detection capabilities against realistic TTPs. It serves as the operational environment for conducting such complex, multi-party exercises.

By developing these more proactive tools, underpinned by advanced cyber capabilities, Australia can move beyond simply reacting to malicious activity and begin to actively shape the cyber environment, making it a more complex, costly, and contested space for our adversaries to operate in.

- ## 3.6 Shield 6: Strong region and global leadership

46. Do you view attributions, advisories and sanctions effective tools for countering growing malicious cyber activity? What other tools of cyber diplomacy and deterrence would you like to see Australia consider for development and use to effectively combat these threats in Horizon 2?

47. Are there additional ways the Australian Government could engage with Southeast Asia or the Pacific to ensure a holistic approach to regional cyber security?

48. Is there additional value that Cyber RAPID can provide in the region beyond its current design and scope?

Cyber RAPID's success in incident response is a tremendous asset, and that trust and brand recognition can be leveraged to provide much deeper, more strategic value beyond its current scope. The key is to evolve it from a purely reactive, tactical tool into a proactive, capability-building program.

Here are three ways Cyber RAPID could be enhanced to provide additional, holistic value in the region.

Enhancement Strategy	Description: How it Would Work	The Additional Value It Provides
1. From Reactive Response to Proactive Resilience	Introduce two new "peacetime" functions under the Cyber RAPID brand: 1. **Cyber RAPID Readiness Assessments:** Small, forward-deployed teams conduct short, confidential assessments with critical infrastructure entities in partner nations. They identify key vulnerabilities and provide a prioritized roadmap for uplift, all under the trusted RAPID banner. 2. **Cyber RAPID Exercise Facilitation:** The teams use their real-world crisis experience to design and facilitate tabletop exercises and simulations for partner governments and companies. They would simulate a realistic	**Prevents Incidents Before They Happen:** This shifts the program's focus from "response" to "prevention and preparedness." It builds muscle memory within the partner nation, making them more self-sufficient. When a real crisis does occur, the RAPID deployment is far more effective because relationships and a baseline understanding already exist.

Enhancement Strategy	Description: How it Would Work	The Additional Value It Provides
	attack and help the local team walk through their response process.	
2. From Technical Fix to Strategic Intelligence	Leverage the unique data and insights gathered from deployments to create a regional intelligence feedback loop: 1. **Regional Threat Briefings:** On a quarterly basis, the Cyber RAPID program would publish an anonymized, aggregated "State of the Threat" report for the region. This report would detail the common adversary TTPs, malware strains, and vulnerabilities observed across all deployments. 2. **Policy & Investment Advisory:** The program would establish a formal channel to provide confidential advice back to both Australian and partner nation policymakers. For example, "We have observed in three separate incidents that a lack of data sovereignty law hindered our investigation," or "The most common point of failure is unpatched firewalls, indicating a need for	**Scales the Benefit to the Entire Region:** The lessons learned from helping one nation can be used to protect all of them. This transforms a series of individual technical engagements into a powerful engine for regional collective defence. **Drives Smarter Policy and Investment:** It provides policymakers with invaluable, evidence-based insights from the front lines. This ensures that future government investments and policy initiatives are targeted at solving the real-world problems that are actually causing harm in the region.

Enhancement Strategy	Description: How it Would Work	The Additional Value It Provides
	investment in asset management."	
3. From "Doing For" to "Training With"	Formalize a capability uplift mandate as a core component of every single RAPID deployment: 1. **Embedded Shadowing:** For every Australian expert deployed, a local counterpart from the affected organization *must* be assigned to shadow them. The primary goal is not just to fix the problem, but to actively transfer knowledge and skills in a real-world crisis environment. 2. **"Train the Trainer" Hotwash:** At the conclusion of a deployment, the RAPID team would conduct a mandatory "hotwash" and training session for the local IT and security teams. This would focus on the specifics of the attack they just faced and train them on how to detect and respond to it themselves in the future.	**Builds Lasting Sovereign Capability:** This is the most crucial evolution. It reframes a deployment as an intensive, hands-on training opportunity, not just a service call. It directly addresses the goal of enhancing regional capabilities by leaving behind a more skilled and confident local workforce, reducing future dependency and cementing Australia's role as a true partner in their long-term success.

By expanding Cyber RAPID in these ways, Australia can transform its most successful tactical tool into a powerful strategic asset. It would evolve from being the region's most trusted *firefighter* to being its most valued *fire marshal*—not just putting out fires, but helping to design a more fire-resistant neighbourhood for everyone.

49. In which forums and on which issues would you like Australia to focus efforts to shape rules, norms and standards in line with its interests most effectively in Horizon 2?

50. What regulatory frameworks or requirements should be prioritised for consideration as part of Australia's efforts on international cyber regulatory alignment?

Pursuing international regulatory alignment is a force multiplier—it can simultaneously reduce the burden on Australian businesses and strengthen our collective defence with key allies. To be effective, Australia's efforts must be focused, pragmatic, and targeted at the areas of greatest friction and highest reward.

We recommend focusing efforts in three key areas, while being mindful of the significant opportunities and challenges involved.

Priority Areas for International Regulatory Alignment

Priority Area	Why it Matters	Key Partners for Alignment
1. Incident Reporting & Threat Sharing	This is the area of highest friction for global companies. Divergent timelines (e.g., 12 hours under SOCI Act, 72 hours under GDPR), thresholds, and reporting formats create immense confusion during a crisis. Harmonizing these requirements would be the single biggest win for industry and for international incident response.	**EU (NIS2 Directive & GDPR), US (CIRCIA), UK, Canada, Singapore.**

Priority Area	Why it Matters	Key Partners for Alignment
2. Critical Infrastructure Protection	As our critical infrastructure is operated by multinational companies and relies on global supply chains, aligning our SOCI Act framework with international equivalents is essential. The goal should be **"compliance reciprocity,"** where demonstrating compliance with a trusted partner's framework (e.g., the EU's NIS2 Directive) is recognized as meeting the majority of Australia's requirements.	**EU (NIS2 Directive), US (CISA Directives), UK (NCSC Cyber Assessment Framework).**
3. Secure Technology & Supply Chain Standards	Australia, along with its partners, is a major consumer of global technology. Instead of creating bespoke Australian standards, we should work with allies to develop a common, high bar for the security of critical technologies (e.g., IoT, managed services, cloud platforms). This creates a powerful, unified market that incentivizes global manufacturers to build security in from the start.	**US (NIST Frameworks), UK, Canada, EU (Cyber Resilience Act).**

Key Opportunities and Challenges

Pursuing this alignment presents immense opportunities, but it must be done with a clear-eyed view of the challenges.

	Opportunities (The Upside)	Challenges (The Hurdles)
For Australian Industry	**Reduced Compliance Burden:** A "comply once, report to many" approach would dramatically lower costs and complexity for Australian businesses operating globally. **Easier Market Access:** Alignment makes it easier for Australian tech companies to sell into allied markets, as they would be building to a common, understood standard.	**Navigating Complexity:** Even with alignment, businesses will still need to navigate the nuances of each jurisdiction's legal system and enforcement regime.
For National Security & Diplomacy	**Enables Machine-Speed Cooperation:** Harmonized reporting formats allow for automated, machine-to-machine threat intelligence sharing between national cyber defence agencies, enabling a faster collective response. **Strengthens Alliances:** The process of working towards alignment deepens diplomatic, economic, and security ties with our most important partners, building a unified front against adversaries.	**National Sovereignty:** Nations are inherently reluctant to cede regulatory control or adopt another country's model wholesale. The key is to aim for alignment of *outcomes*, not identical legislative text. **Pace of Change:** The legislative process is slow, while the threat landscape evolves rapidly. Alignment efforts must be agile, focusing on high-level principles and leveraging regulatory dialogues rather than slow-moving international treaties.

	Opportunities (The Upside)	Challenges (The Hurdles)
For the Cyber Ecosystem	**Raises the Global Security Baseline:** By creating a large, aligned bloc of major economies, we can exert significant influence on global technology manufacturers to prioritize security, benefiting the entire global ecosystem.	**Divergent Legal Traditions:** Deep-seated differences between legal systems (e.g., common law vs. civil law, different approaches to privacy and government powers) can create complex and difficult-to-resolve barriers to full alignment.

In conclusion, pursuing regulatory alignment is the logical next step in Australia's cyber security maturation. By focusing on high-friction areas and working pragmatically with key allies, we can reduce the burden on our businesses and build a more seamless and effective collective defence for the future.

And the question is -

What regulatory frameworks or requirements should be prioritised for consideration as part of Australia's efforts on international cyber regulatory alignment?

This is an excellent, focused question. To maximize impact with limited resources, Australia should prioritize aligning on frameworks that address the most immediate cross-border frictions and offer the greatest mutual security benefits.

We recommend a tiered approach, focusing first on the operational realities of incident response, then on the strategic protection of critical infrastructure, and finally on the long-term goal of securing the future technology supply chain.

Proposed Priorities for International Cyber Regulatory Alignment

Priority Tier	Regulatory Area	Key International Counterparts & Frameworks	Rationale for Prioritization (The "Why")
Priority 1: Urgent	**Incident & Data Breach Reporting**	**EU:** GDPR & NIS2 Directive **US:** CIRCIA	**Addresses the Highest Friction:** This

Priority Tier	Regulatory Area	Key International Counterparts & Frameworks	Rationale for Prioritization (The "Why")
		(Cyber Incident Reporting for Critical Infrastructure Act) **UK:** UK GDPR & NIS Regulations **Singapore:** CSA's licensing framework	is the most acute pain point for any organization operating internationally. During a crisis, navigating a patchwork of different reporting timelines (e.g., 12, 24, 72 hours), thresholds, and data formats is a significant and costly burden that slows down response for both the company and national security agencies. **Enables a Common Operating Picture:** Alignment on *what* constitutes a reportable incident and *what information* needs to be shared would allow for the creation of a common, machine-readable format. This would enable automated threat sharing between allied national cyber security centres, dramatically accelerating collective defence.

Priority Tier	Regulatory Area	Key International Counterparts & Frameworks	Rationale for Prioritization (The "Why")
Priority 2: Strategic	**Critical Infrastructure Protection Frameworks**	**EU:** NIS2 Directive (Network and Information Security) **US:** CISA Directives & NIST Cybersecurity Framework (CSF) **UK:** NCSC Cyber Assessment Framework (CAF)	**Drives "Compliance Reciprocity":** The core goal is to establish a system where demonstrating compliance with a trusted partner's framework (e.g., NIS2) is accepted as meeting the vast majority of the requirements under Australia's SOCI Act (specifically the CIRMP). This would massively reduce the audit and compliance burden for the multinational companies that operate our critical infrastructure. **Strengthens Shared Resilience:** As critical infrastructure increasingly relies on a global supply chain, aligning our core protection frameworks ensures a consistently high bar of security across all allied

Priority Tier	Regulatory Area	Key International Counterparts & Frameworks	Rationale for Prioritization (The "Why")
			nations, making the entire system more resilient to attack.
Priority 3: Forward-Looking	**Secure Technology & Supply Chain Standards**	**EU:** Cyber Resilience Act (CRA) & RED (Radio Equipment Directive) **US:** NIST's Secure Software Development Framework (SSDF) & IoT security guidance.	**Shapes the Global Market:** Rather than creating a unique "Australian standard," aligning with our partners on security requirements for hardware and software products (especially IoT) creates a powerful, unified market. This incentivizes global manufacturers to build secure-by-design products for everyone, raising the security baseline globally. **Reduces Systemic Risk:** By agreeing on a common framework for vetting technology suppliers and assessing software supply chain security, we can collectively reduce the risk of a

Priority Tier	Regulatory Area	Key International Counterparts & Frameworks	Rationale for Prioritization (The "Why")
			systemic vulnerability (like SolarWinds) impacting all our economies simultaneously.

Recommendation

Australia should initiate formal "Regulatory Alignment Dialogues" with the EU, US, and UK, with Singapore as a key regional partner. The first order of business in these dialogues should be to map and seek alignment on **incident reporting requirements**. Achieving even partial success in this area would deliver an immediate and significant benefit to both Australian industry and our national security posture.