Board Cyber-AI Risk Questions

Board and executive questions on cyber and AI risk, mapped to CPS 234, CPS 230 and s912A

cydarm

www.cydarm.com

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

Background

In the space of nine days, both of Australia's financial regulators wrote to the entities they oversee about artificial intelligence.

30 April 2026 – APRA, to all regulated entities: Observations and expectations on artificial intelligence. The letter signals that existing obligations under CPS 234 and CPS 230 now apply to the AI era, and that boards must provide genuine challenge, not simply receive and accept management briefings. APRA has indicated that enforcement action will follow if entities fail to manage these risks.

8 May 2026 – ASIC, to AFS licensees and market participants: Parallel expectations for AFS licensees, issued as an open letter, anchored in the FIIG Securities precedent. For dual-regulated entities (banks with wealth or broking arms, super trustees, insurers with advice businesses) the same standard now applies under s912A of the Corporations Act, not just under CPS 234 and CPS 230.

This paper is a direct response to both letters.

Regulatory Grounding

CPS 234 requires regulated entities to maintain information security capability "commensurate with the size and extent of threats" - a standard that flexes as the threat environment changes. ASIC's parallel formulation, articulated in ASIC v FIIG Securities Limited, requires controls to be "demonstrably effective and proportionate to the size, nature and complexity of the business."

The questions in this paper are grounded in those two standards and in the prudential framework that supports them - CPS 234, CPS 230 and CPS 220, together with s912A of the Corporations Act for AFS licensees. They are intended to give board members the tools to seek meaningful assurance, not to turn directors into technical experts. This paper is written for APRA-regulated entities, including those that also hold an Australian Financial Services Licence. For dual-regulated entities, the AFSL references apply in addition to APRA prudential obligations; for APRA-only entities, those references are contextual.

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

What Has Changed

Two conditions have materially changed the context in which these questions should be asked.

The first is external. Artificial intelligence has significantly accelerated the pace at which adversaries can identify vulnerabilities, develop exploits, and launch attacks. The window between a vulnerability being discovered and it being actively exploited has narrowed considerably. Threats that previously took weeks to materialise can now emerge in hours. This means the speed and resilience of your organisation's response capacity matters more than it ever has.

The second is internal. AI tools are being adopted rapidly across financial services, in customer service, software development, document processing, and operations. Each adoption changes your organisation's risk surface. The internal risk surface includes AI use that has been formally adopted by the organisation as well as AI tools being used by staff outside approved control frameworks. The question is not whether to use AI, but whether the risks created by its use are being identified, assessed, and managed with the same rigour that APRA expects for any other material operational risk.

Together, these conditions mean that a cyber posture that was adequate twelve months ago may not be adequate today. The questions that follow are designed with that reality in mind.

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

1. Know Your Exposure

AI adoption inside the business is changing the risk surface continuously. The questions in this section are not just about understanding current exposure; they are about whether the organisation has the visibility to keep pace with its own rate of change.

1. Do we have a current inventory of our most critical information assets, and is ownership clearly assigned? (CPS 234, paras 13, 15 & 20; AFSL: s912A(1)(a), (h))
2. What is the full extent of our known security weaknesses across the entire estate, and who is accountable for each one? (CPS 234, paras 14, 15 & 36; AFSL: s912A(1)(h))
3. How do we manage the risk that our own people, whether through error or intent, could compromise our security? (CPS 234, para 13; AFSL: s912A(1)(d), (h))
4. How confident are we in the security practices of our most critical suppliers? (CPS 234, paras 16, 21 & 28; CPS 230, paras 26–38; AFSL: s912A(1)(h))
5. How dependent are we on a small number of critical technology providers, and what happens if one fails or is compromised? (CPS 230, paras 26–38; AFSL: s912A(1)(h))
6. Do we know what software components are running in our systems, including those we didn't build? (CPS 234, para 15; AFSL: s912A(1)(h))
7. Where has AI been introduced into our operations, and what risks has that created? (CPS 230, para 17; AFSL: s912A(1)(h))
8. How do we know what AI tools our people are using that we haven't formally approved, and what data are those tools touching? (CPS 234, paras 13 & 15; CPS 230, para 17; AFSL: s912A(1)(h))
9. What controls are in place to prevent unapproved AI use on sensitive data? (CPS 234, para 13; AFSL: s912A(1)(a), (h))

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

10. How are we ensuring that the way we use AI internally does not create risks faster than our security posture can adapt to them? (CPS 230, para 17; AFSL: s912A(1)(h))

11. Are our controls on unsanctioned AI use preventative and technically enforced, or do we rely primarily on policy and after-the-fact detection? (CPS 234, paras 13 & 15; AFSL: s912A(1)(h))

12. How do we test the AI systems we rely on for weaknesses that are specific to AI, and are those tests independent of the teams that built or deployed them? (CPS 234, para 16; AFSL: s912A(1)(h))

13. Do we have the right tools to detect vulnerabilities across our entire estate, and how do we validate that those tools are actually capturing what they should? (CPS 234, para 15; AFSL: s912A(1)(h))

14. Do our development teams use open source software, and if so, how is that usage governed and tracked? (CPS 234, para 15; AFSL: s912A(1)(h))

15. Can we trace the origin and review history of the software components in our critical systems, and would we know if something had changed since it was last reviewed? (CPS 234, para 15; CPS 230, para 17; AFSL: s912A(1)(h))

16. For our critical AI and technology suppliers, do our contracts provide for audit rights, advance notification of model or data handling changes, incident triggers, and tested exit or substitution plans, and have we aligned material AI providers to our CPS 230 material service provider framework? (CPS 230, paras 26–38; AFSL: s912A(1)(h))

17. For AI systems that influence high-impact decisions, where is human involvement required, and how do we evidence that the human role is substantive rather than procedural? (CPS 230, para 17; CPS 234, para 13; AFSL: s912A(1)(a), (h))

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

2. Reduce & Contain Exposure

18. To what extent does our current security posture rely on perimeter controls to manage vulnerabilities? Has our approach adapted to attackers using AI tools to identify and exploit vulnerabilities, including chaining multiple low impact vulnerabilities for greater impact? (CPS 234, paras 15 & 16; AFSL: s912A(1)(d), (h))
19. What active steps are we taking to reduce the exposure of our systems and services to threats, and how is that progress measured? (CPS 234, para 15; AFSL: s912A(1)(d), (h))
20. Is our security architecture designed on the assumption that we will be breached, and how do we limit impact when that happens? (CPS 234, paras 15 & 16; AFSL: s912A(1)(d), (h))
21. Where are we using AI defensively - to identify vulnerabilities, review code, or detect and respond to threats - and how do we assess whether that use is effective? (CPS 234, paras 15 & 16; AFSL: s912A(1)(d), (h))

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

3. Speed & Resilience

AI has compressed the time between a vulnerability emerging and it being exploited. Organisations that measure their response capacity in weeks are operating in a threat environment that now moves in hours. The questions in this section probe whether our response capability is calibrated to that reality.

22. How quickly can we respond to a newly discovered vulnerability, and how do we know? (CPS 234, para 15; AFSL: s912A(1)(d), (h))
23. What would it take to patch every system we operate, and how long would that take today? (CPS 234, para 15; AFSL: s912A(1)(d), (h))
24. When something goes wrong, how long before we know about it? (CPS 234, para 20; AFSL: s912A(1)(d), (h))
25. If we suffered a serious breach today, what is our plan and how recently was it tested? (CPS 234, para 20; AFSL: s912A(1)(d), (h))
26. How long would it take to restore our critical services after a cyber incident? (CPS 230, para 20; AFSL: s912A(1)(d), (h))
27. How is our ability to detect and respond to threats keeping pace with the speed at which those threats are evolving, and what evidence do we have of that? (CPS 234, paras 15 & 20; AFSL: s912A(1)(d), (h))
28. Beyond patching, how much of our response to a cyber incident relies on automated detection and action versus human decision-making, and is that balance appropriate given the speed of today's threats? (CPS 234, paras 15 & 20; AFSL: s912A(1)(d), (h))
29. If we are now patching at the cadence the current threat environment requires, how are we governing the risk that the patching process itself introduces - identification, testing, and oversight of changes happening at higher velocity? (CPS 234, para 15; AFSL: s912A(1)(d), (h))

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

4. Assurance & Change Control

30. Who tests our defences independently, how often, and what have they found? (CPS 234, para 16; AFSL: s912A(1)(a), (h))
31. How do we know our security measures are actually working, not just in place? (CPS 234, para 16; AFSL: s912A(1)(a), (h))
32. What role does internal audit play in independently assessing our cyber posture, and is it adequately resourced? (CPS 234, para 17; AFSL: s912A(1)(a), (h))
33. How do we ensure that changes to our systems and processes don't introduce new security risks? (CPS 234, para 15; AFSL: s912A(1)(h))
34. What are our obligations to notify APRA when a security incident occurs, and do we have a process that would meet those obligations under pressure? (CPS 234, paras 18–19; AFSL: s912A(1)(h))
35. How do we validate the accuracy of our security metrics, and who checks the checkers? (CPS 234, paras 16 & 17; AFSL: s912A(1)(a), (h))
36. For AI systems and other probabilistic technologies, how are we moving beyond point-in-time and sample-based assurance toward continuous validation of behaviour, drift and control effectiveness? (CPS 234, paras 16 & 17; AFSL: s912A(1)(a), (h))

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

5. Board Oversight

37. What is our position on how much cyber risk we are willing to accept, and are we staying within it? (CPS 220, para 24; AFSL: s912A(1)(a), (h))

38. How are cyber risk settings reviewed and updated as the threat environment changes? (CPS 234, para 12; AFSL: s912A(1)(a), (h))

39. What would have to happen before this board would conclude that our cyber framework is no longer adequate? (CPS 234, para 12; AFSL: s912A(1)(a), (h); RG 104)

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com

6. Executive Governance & Accountability

40. How does management demonstrate that it is operating within the risk appetite the board has set? (CPS 220, para 24; AFSL: s912A(1)(a), (d); RG 105)

41. What does management escalate to the board, and what does it handle itself, and is that the right division? (CPS 234, para 12; AFSL: s912A(1)(a), (d); RG 105)

42. What is the audit committee doing to satisfy itself, not just to receive reports? (CPS 234, para 12; AFSL: s912A(1)(a), (d); RG 105)

This paper does not constitute legal or regulatory advice. Regulated entities should seek independent advice on their obligations under CPS 234, CPS 230, CPS 220, and related APRA standards, and (for AFS licensees) s912A of the Corporations Act and related ASIC guidance including RG 104 and RG 105.

Cydarm Technologies
121 King St, Melbourne, VIC
info@cydarm.com