AI GOVERNANCE · FRONTIER TECHNOLOGY · LEADERSHIP

Mind the Gap

Why an AI governance framework is not enough for your organisation

Bernadette Harkin GAICD | Founder and Principal, Lucentis Advisory

If you have ever stood on a platform waiting for a train, you know the phrase. It is a warning about the space between where you are and where you need to be. This paper is about a different gap, one I discuss almost every day with boards and executive teams: the space between choosing an AI governance framework and actually governing AI.

Choosing a framework should not feel like picking a side. But for a lot of boards right now, that is exactly what it feels like. There is the National AI Centre's guidance, OECD principles, APRA's technology risk lens, and a growing list of industry body recommendations, just for a start. And that is before you get to the sector-specific standards and the international instruments layered on top.

Each one is credible. Each one is asking for a board's attention. And none of them quite fit together out of the box.

The question I hear most often from directors is not "should we govern our AI?" Most boards have moved past that question. The one they are actually asking is "which framework do we use?" And that question, on its own, is a trap. It assumes the answer to AI governance is a document you can select off a shelf. It is not.

A crowded landscape, and why that is the trap

Start with what is actually on the table. The National AI Centre, Australia's government body for AI policy and adoption, publishes guidance aimed at helping organisations use AI safely and responsibly. The OECD has its own AI Principles, adopted by dozens of member and partner countries, setting out high-level commitments around transparency, accountability and human rights.

APRA brings a technology and operational risk lens that applies directly to regulated financial entities and flows indirectly into how their suppliers and partners are expected to operate. Layer on sector-specific standards — whatever applies to healthcare, professional services, critical infrastructure or education — and then the international instruments that shape the direction of travel for regulation more broadly.

Each of these was written by a different body, for a different purpose, at a different point in time. None of them were written with your organisation, its purpose, its risk appetite or its

stakeholders, specifically in mind. That is not a criticism of any single framework. It is simply what happens when governance guidance is produced in parallel by bodies with different mandates.

For a board, the practical effect is that adopting any one framework wholesale still leaves gaps that the others were designed to cover, and reconciling all of them manually is not a realistic use of anyone's time. So when a director asks which framework to use, the honest answer is that the choice of framework was never the decision that mattered most. The decision that matters is what your organisation does with any of them once chosen.

What the International AI Safety Report actually tells boards

It helps to be specific about the evidence base boards are working from. The International AI Safety Report is now in its second edition, published in February 2026, chaired by Turing Award winner Yoshua Bengio. Authored by more than 100 AI experts, backed by over 30 countries and international organisations including the OECD and the UN. It grew out of the mandate world leaders set at the 2023 AI Safety Summit at Bletchley Park.

The Report frames general-purpose AI risk in three categories: malicious use, malfunctions, and systemic risks. That framing matters because it moves the conversation beyond a single worry — such as data privacy or bias — into a fuller picture of how these systems can fail or be misused, and how the consequences can ripple through an organisation and the systems it depends on.

On capability, the Report is careful about uncertainty, but its trend extrapolations suggest that models could reach research-level performance across specialised scientific domains within the next few years. Frontier models achieved gold-medal-level performance at the International Mathematical Olympiad in 2025. None of that is theoretical for a board. It means the competitive pressure to deploy is real, and it is not going away while a governance framework is debated in committee.

The Report itself acknowledges that the inner workings of frontier models remain poorly understood, and that the relationship between benchmark performance and real-world performance is not yet well mapped. That is precisely why traditional risk management practices, built for systems whose behaviour could be fully specified in advance, were not built to catch what is coming. A framework alone does not close that gap. Interpretation does.

The three lenses

If the answer is not a framework off the shelf, where does that leave a mid-size board? It leaves you needing to hold three things at once — and they are already fundamental to how your organisation operates.

Your organisational purpose. Stakeholder expectations. The direction of travel of regulation. Read every applicable framework through those three lenses, and a generic document becomes something that actually fits your organisation.

The first lens is purpose: your organisation's core values, mission and reason for being. No external framework can define this for you, and reducing a governance decision down to what a particular AI model can technically do skips over the one question only your board can answer — which is whether a given use of AI actually serves what your organisation is here to do.

The second lens is stakeholder expectations: what your investors, customers, employees and communities are watching for. These expectations are often more demanding, and more specific to your context, than anything written into a general framework, and they shift faster than any document can be updated.

The third lens is the regulatory direction of travel: not just the rules as they stand today, but where the guidance from bodies like the National AI Centre, the OECD and APRA is heading. Reading that trajectory matters more than reading any single instrument in isolation, because it tells you where the compliance bar is moving, not just where it sits right now.

The work is not choosing between the NAIC's guidance, the OECD principles, APRA's technology risk lens or a sector standard. It is taking whichever of those genuinely applies to your organisation and reading it through purpose, stakeholders and regulatory direction, together. That is what turns a generic framework into something your organisation can actually govern by.

What good looks like

Done well, this is not compliance with one document. It is interpretation, applied consistently across every framework that touches your organisation, and it produces something genuinely valuable: a balanced, restrained approach to AI deployment. Not paralysis, where a board sits on its hands waiting for perfect clarity that will never arrive. And not recklessness, where a board adopts AI at the pace of a competitor without understanding what it is accountable for.

This is why professional support in assessment and risk analysis matters so much right now — not to hand a board a framework it could have found itself, but to help it do the interpretation that turns frameworks into governance. And this capability needs to be built here in Australia specifically, so that our organisations maintain their competitive advantage and their own organisational culture, rather than importing a version of governance built for a different market and a different set of stakeholders.

Mind it

Legislators around the world are continuing vital work, trying to make these models more predictable, observable and accountable. That work needs to continue. But boards need to be honest about where they are today: on a journey that has already left the station.

Until I find a better way to put it: the frameworks are not the gap. The gap is what happens after you have read them. Mind it.

REFERENCES

International AI Safety Report (2026). Chaired by Yoshua Bengio, authored by over 100 AI experts, backed by more than 30 countries and international organisations. Second edition, published 3 February 2026. Available at internationalaisafetyreport.org/publication/international-ai-safety-report-2026

National AI Centre. Australian Government, Department of Industry, Science and Resources. Available at industry.gov.au/science-technology-and-innovation/technology/artificial-intelligence/national-ai-centre

ABOUT THE AUTHOR

Bernadette Harkin GAICD is the Founder and Principal of Lucentis Advisory, an AI governance consultancy advising boards and executive leadership teams at mid-size organisations. She is also a Non-Executive Director at Archer Materials, Australia's only listed quantum technology company. A career technology executive, she brings decades of experience across technology, consulting and sustainable business growth, built across global markets.

Speaking and media enquiries: bernadette@lucentisadvisory.com.au

