

AI REGULATION IN THE BOARDROOM

What boards must understand about exposure,
accountability and readiness

MANOJ TAVARAJOO



Executive Summary

AI regulation is no longer assessing what organisations intend to do. It is assessing what they can prove.

This shift changes the role of the board.

AI regulation is accelerating globally, yet organisational readiness remains limited. Casepoint's March 2026 analysis found that only 3% of organisations are fully prepared for current AI regulatory requirements, while 52% have minimal or no structured readiness. A separate April 2026 readiness review of fifty European AI companies found that 74% trigger high-risk classification under the EU AI Act, and 96% have no public statement of their position on the regulation. The gap between regulatory expectation and organisational capability is not closing. It is widening.

This paper does not interpret legal obligations. It focuses on how boards and senior executives should respond to regulatory expectations through governance, accountability, evidence and assurance.

Across major regulatory frameworks, there is convergence around a consistent set of expectations. These include risk-based classification, transparency and documentation, human oversight, lifecycle monitoring and clear accountability for outcomes. At the same time, regulatory approaches remain fragmented and continue to evolve, particularly in areas such as general-purpose AI and foundation models.

In Australia and similar jurisdictions, the direction is increasingly principles-led, with regulators emphasising accountability, risk management and operational resilience rather than prescriptive rules. This reinforces the need for organisations to build internal capability rather than rely on compliance checklists.

The implication for boards is not legal complexity. It is operational accountability.

Policies are no longer sufficient. Frameworks are no longer sufficient. Organisations must demonstrate that AI systems are operating within defined boundaries in practice.

**Evidence is the
currency of regulatory
confidence.**

The central question for boards is no longer whether the organisation understands the regulation. It is whether the organisation can operate in a way that meets regulatory expectations consistently and at scale.


Why AI Regulation Is Now a Board Issue

AI regulation has moved beyond legal and compliance functions and into board-level accountability.

This is not simply a continuation of previous technology regulation. It represents a shift in how regulators assess organisational behaviour.

In earlier waves of regulation, organisations were required to demonstrate that appropriate policies, controls and governance structures were in place. In the context of AI, regulators are increasingly focused on whether those controls operate effectively in practice and whether outcomes remain within acceptable boundaries.

This distinction is critical. Organisations that rely on policy-led compliance will find themselves exposed. Regulatory scrutiny is now outcome-based and evidence-driven.



At the same time, readiness across organisations remains uneven. Many organisations are still at an early stage of translating regulatory expectations into operational capability. This creates a widening gap between what regulators expect and what organisations can demonstrate.

Boards are therefore expected to oversee not only whether governance exists, but whether it functions. **They must understand:**

Where AI is being used and for what purpose

What level of risk it introduces

Whether governance and controls are proportionate

Whether there is sufficient evidence and assurance

This is not about understanding legal detail. It is about ensuring that the organisation can operate within regulatory expectations and demonstrate that this is happening.

In Australia, this dynamic is already visible. The Australian Prudential Regulation Authority has signalled an increasing focus on AI governance as part of its operational risk supervisory programme, and the gap between the pace of AI adoption and the maturity of assurance capability has been identified as a specific area of concern. Where AI-specific legislation is still developing, sector regulators are stepping in. This pattern is not unique to Australia. It reflects a broader global reality in which the absence of a single AI law does not mean the absence of regulatory expectation. The expectations exist. They are being enforced through existing regulatory powers, and boards are accountable for whether the organisation can meet them.

The Direction of AI Regulation

The most important feature of AI regulation is not the specific content of individual laws. It is the convergence of expectations around how organisations must operate.

Across the EU AI Act, NIST AI Risk Management Framework, ISO/IEC 42001 and OECD principles, there is increasing alignment around a common operating model for responsible AI. **This operating model includes:**

Classification of systems based on risk

Defined requirements for higher-risk use cases

Transparency and documentation

Human oversight

Ongoing monitoring across the lifecycle


Clear accountability for outcomes

This convergence matters more than jurisdictional detail. It means organisations are not responding to isolated regulations. They are responding to a consistent set of expectations about how AI should be governed and operated.

At the same time, regulatory uncertainty remains a defining feature. Areas such as general-purpose AI, foundation models and cross-border accountability continue to evolve. This uncertainty is not temporary. It is structural.

Organisations must therefore build governance models that can absorb change, not just comply with current rules.

The shape of that uncertainty differs by jurisdiction but the governance implication is consistent. In the European Union, implementation timelines have shifted and the treatment of general-purpose AI models continues to be refined. In the United States, regulation remains fragmented across federal agencies and state legislatures, with Colorado, California and Texas among the states advancing their own AI-specific frameworks alongside sector regulators applying existing powers. In Australia, the government has moved away from a dedicated AI law in favour of a principles-based approach, with APRA, ASIC and the OAIC applying existing regulatory mandates to AI-related risks. Boards should not treat this uncertainty as a reason to defer action. Regulatory uncertainty must be governed, not waited out. The governance capability required to operate responsibly under current expectations is the same capability that will be required as those expectations develop.



What Boards Need to Understand

AI regulation is often framed as a compliance obligation. For boards, this framing is incomplete.

The primary implication is that organisations must operate in a way that is observable, controllable and defensible.

Regulatory exposure is created by how AI is used. Systems that influence customer outcomes, employee decisions or financial results carry different levels of exposure depending on their impact and context.

Exposure is not uniform. Within a single organisation, some AI systems may carry minimal regulatory risk, while others may fall into higher-risk categories. This requires differentiated governance.

For example, an internal model used for process optimisation may require limited oversight, while a customer-facing decision system affecting eligibility or pricing may require significantly stronger controls, monitoring and accountability.

Regulators are increasingly focused on evidence. It is no longer sufficient to define policies or governance frameworks. Organisations must demonstrate that controls are in place and operating effectively.

Evidence is the currency of regulatory confidence.

Regulatory expectations extend across the lifecycle of AI systems. Governance must include monitoring, performance management and response to change over time.

Accountability must be clear and operational. Boards should expect defined ownership supported by evidence that accountability is functioning in practice.

Boards do not need to interpret regulation. They need to ensure that the organisation can operate within it and demonstrate that this is happening.

Board Oversight Boundaries

Boards are most effective when they focus on governance, not execution.

AI regulation reinforces this principle. It requires boards to define direction, set risk appetite and ensure accountability, rather than becoming involved in control design or compliance execution. This boundary is not a limitation. It is what makes governance effective. Boards create value by defining acceptable boundaries for AI use, ensuring governance structures are in place, confirming that accountability is clearly assigned, requiring evidence that controls are functioning, and intervening when risk exceeds acceptable thresholds.

Management is responsible for interpreting regulatory requirements, designing controls, operating systems and managing compliance on a day-to-day basis. This distinction is critical. If boards move into execution, they dilute their effectiveness and blur accountability. If they remain focused on governance, they strengthen oversight and ensure that regulatory expectations are met through disciplined organisational capability.



Regulatory Exposure Across the AI Lifecycle



Regulatory exposure is inherently multi-dimensional. It is created at the intersection of use case, regulatory domain, risk level and jurisdiction.

AI systems do not exist within a single regulatory framework. They intersect with privacy, cyber security, consumer protection and sector-specific regulation. Exposure also varies by use case. A system influencing customer outcomes carries a different risk profile from an internal operational tool. The level of impact determines the level of scrutiny.

Jurisdiction adds further complexity, particularly for organisations operating across regions. A single AI system deployed across multiple markets may simultaneously engage EU AI Act obligations, US state-level requirements, Australian prudential expectations and sector-specific rules. Each domain applies with different thresholds and evidence expectations.




Figure 1: AI Regulatory Exposure Map — An enterprise view of where AI creates regulatory exposure

The AI Regulatory Exposure Map provides a structured way to understand this complexity. For boards, it enables identification of high-exposure areas, prioritisation of governance effort and more informed oversight discussions.

Without this clarity, organisations risk misallocating effort, applying insufficient control in high-risk areas and excessive control in low-risk areas. The map is a working tool, not a compliance checklist. It should be revisited as use cases evolve, regulatory frameworks develop and the organisation's AI footprint grows.



The Regulation to Operating Model Framework

The central challenge in AI regulation is not interpretation. It is translation.

Regulation defines expectations, but it does not prescribe how organisations should operationalise them. This creates a gap between regulatory intent and operational reality that policies and frameworks alone cannot close.

Closing this gap requires a structured approach. First, regulatory requirements must be interpreted in context. Concepts such as risk classification, transparency and oversight must be mapped to how AI is used within the organisation. Second, this interpretation must be embedded in the operating model, including roles, responsibilities, decision rights and governance structures. Third, controls and processes must be designed to ensure that AI systems operate within defined boundaries. Fourth, organisations must establish mechanisms to generate evidence through monitoring, reporting and documentation. Fifth, assurance must be in place, with independent validation confirming that governance and controls are functioning as intended.

The mechanism that makes this translation operational is the behavioural envelope: the defined set of limits within which each AI system is expected to operate. The envelope captures dimensions including accuracy, bias thresholds, latency, cost, drift over time and degree of permitted autonomy. Each dimension is linked to a measurement approach, a defined threshold and an escalation trigger. In regulatory terms, the behavioural envelope is how risk classification becomes an operational boundary, how human oversight requirements become intervention thresholds, how monitoring obligations become continuous measurement, and how incident reporting requirements become defined escalation paths. Boards approve the envelope. Management operates within it. Evidence demonstrates that it is holding.

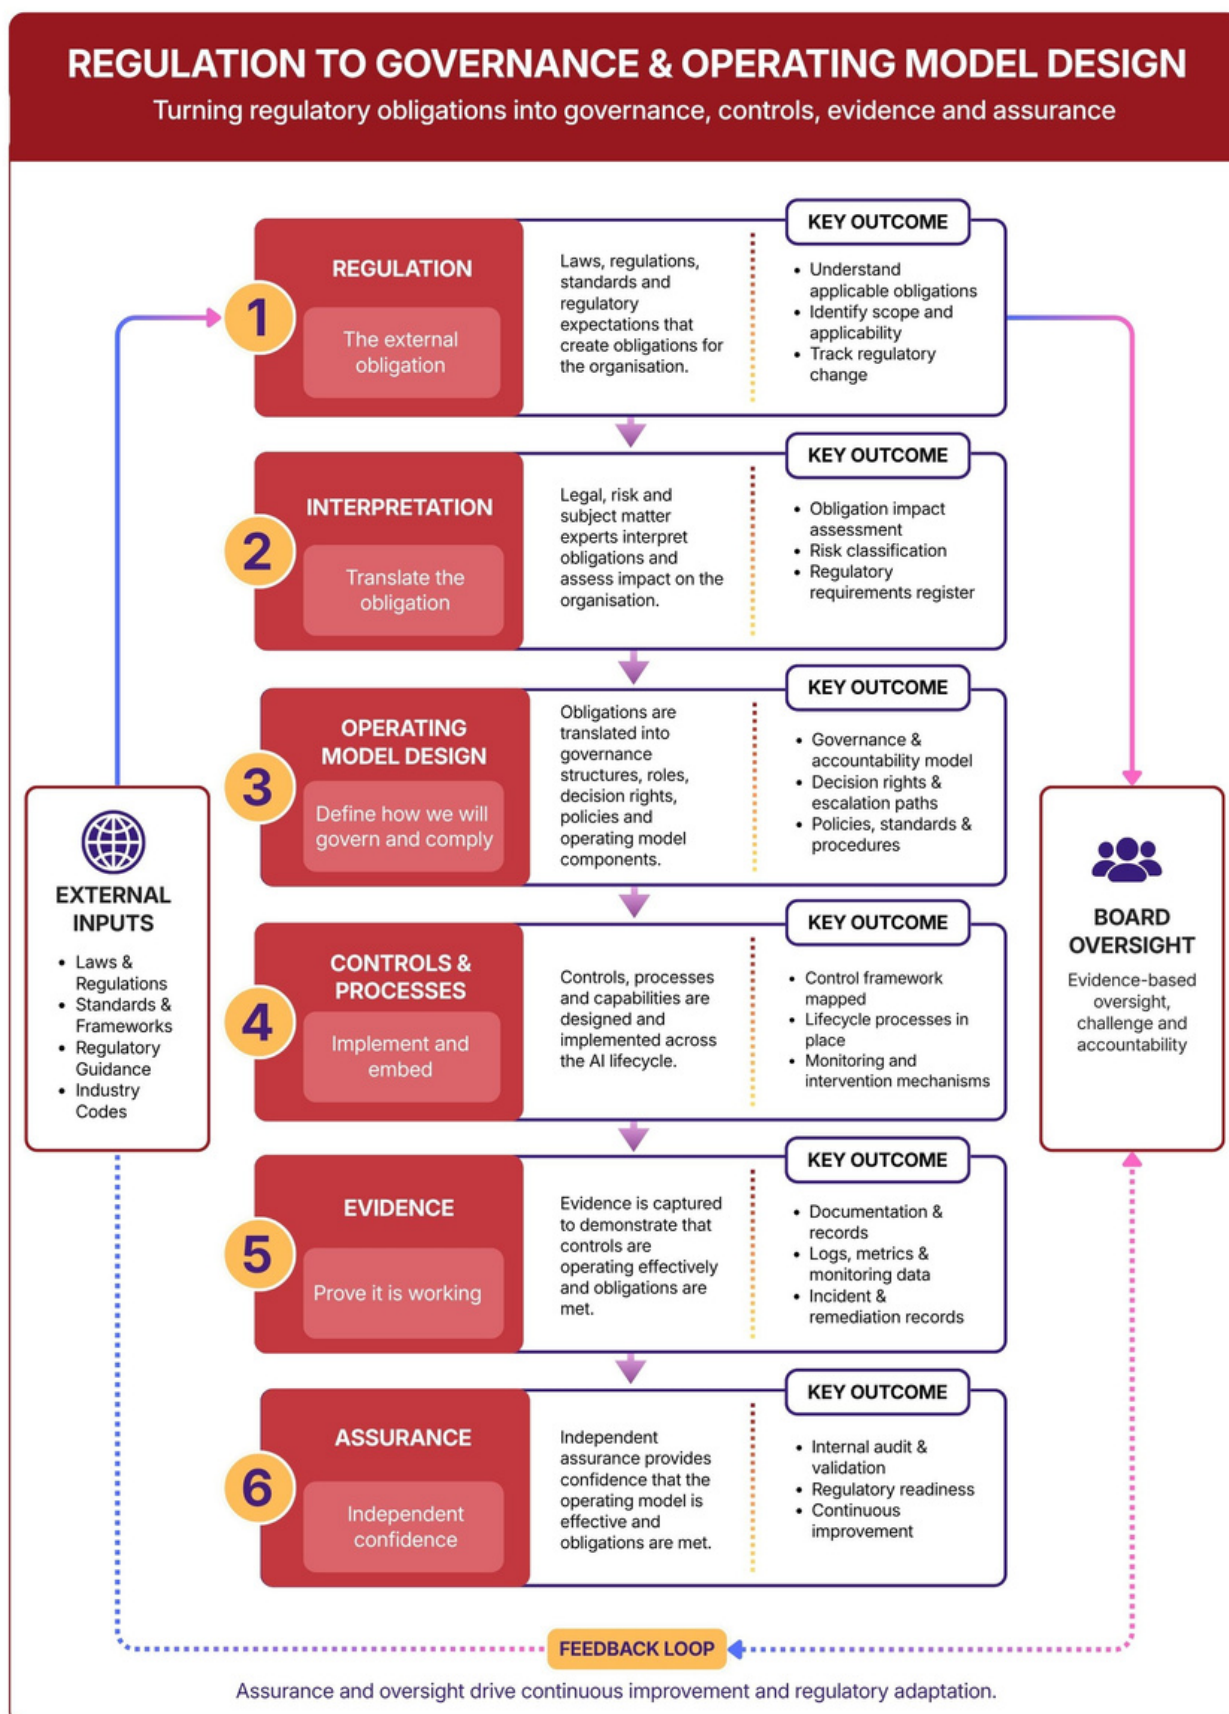


Figure 2: Regulation to Governance & Operating Model Design

This is not a linear process. It is a continuous adaptive system.

Regulatory expectations evolve. Organisational use of AI evolves. New risks emerge. Existing controls degrade over time. The operating model must continuously respond to these changes.

Assurance and oversight generate insight that feeds back into governance, operating model design, controls and monitoring. Regulatory change similarly feeds back into interpretation and design. This creates a closed-loop system in which regulation informs design, design drives execution, execution generates evidence, evidence enables assurance, and assurance drives improvement.

Regulatory readiness is not a one-time exercise. It is an ongoing capability.

Management Questions Boards Should Ask

Boards govern AI regulation through structured, evidence-based inquiry. The questions below are organised around key areas of regulatory oversight and are specifically calibrated to regulatory exposure, evidence of compliance and readiness under scrutiny.

Figure 3 provides a structured question framework that boards can use to organise their regulatory oversight inquiry across the five key areas of exposure, classification, evidence, accountability and readiness.



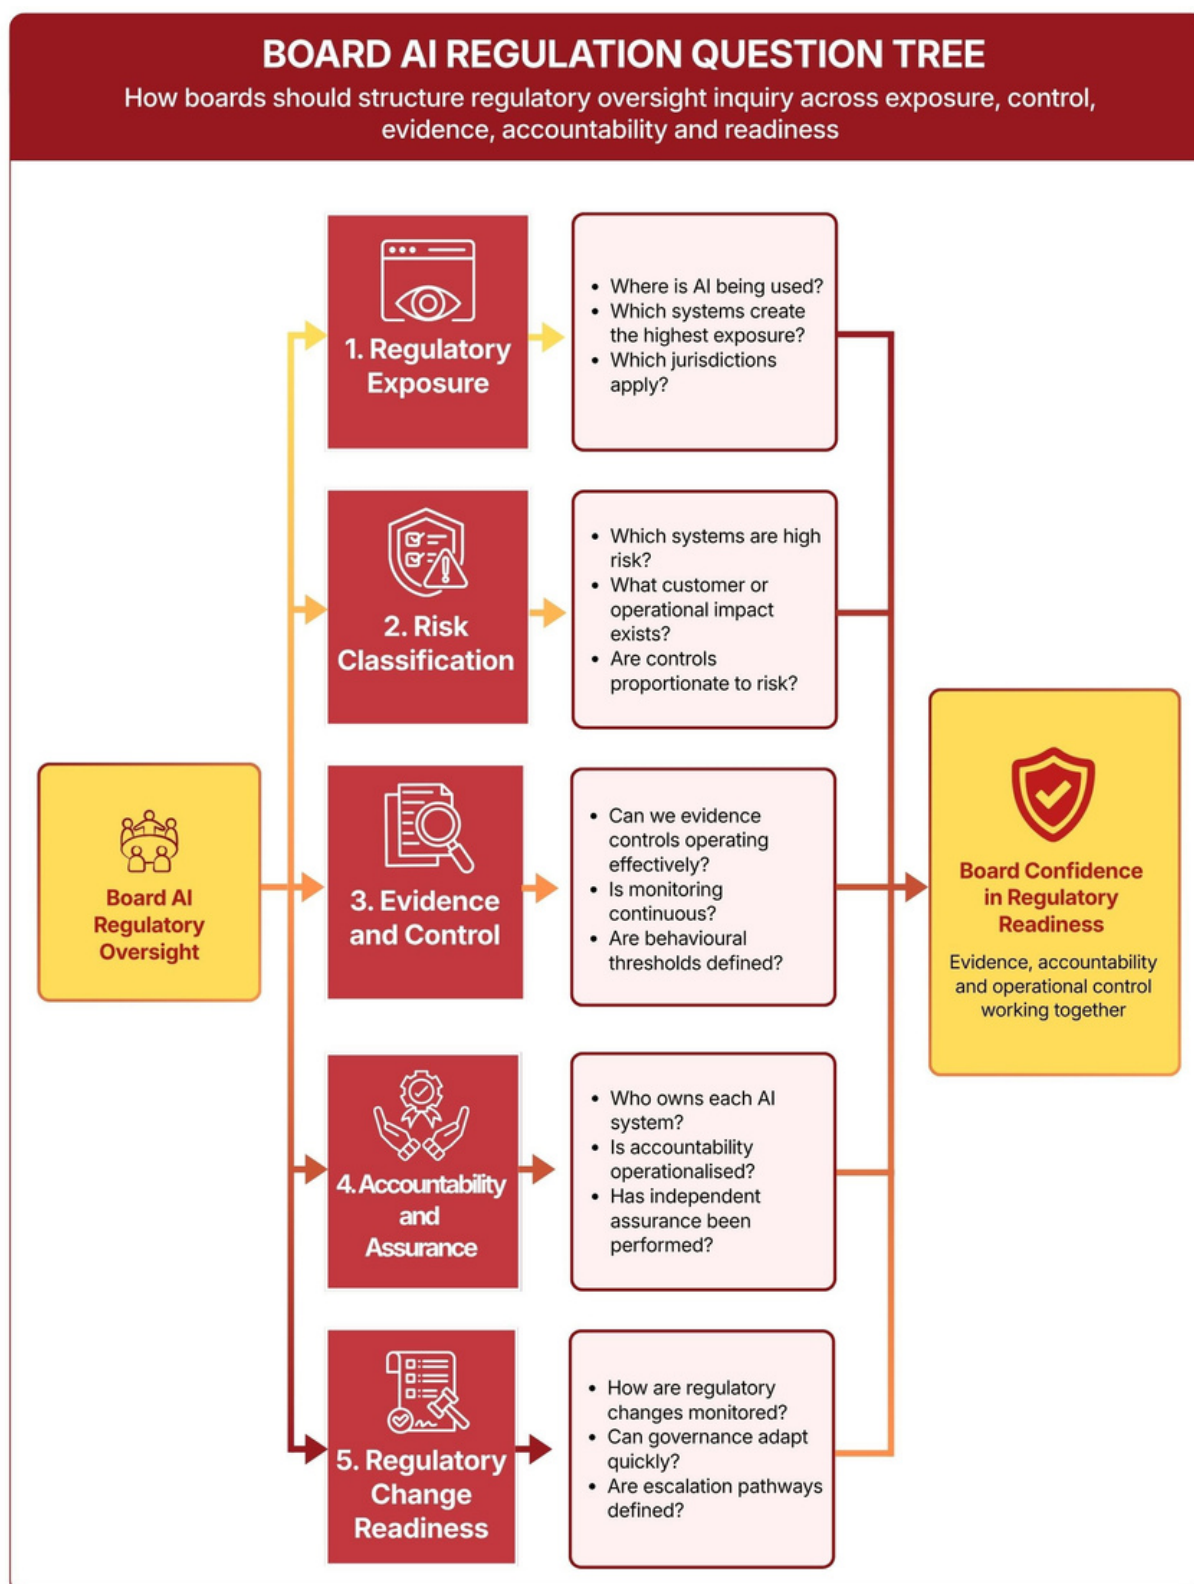
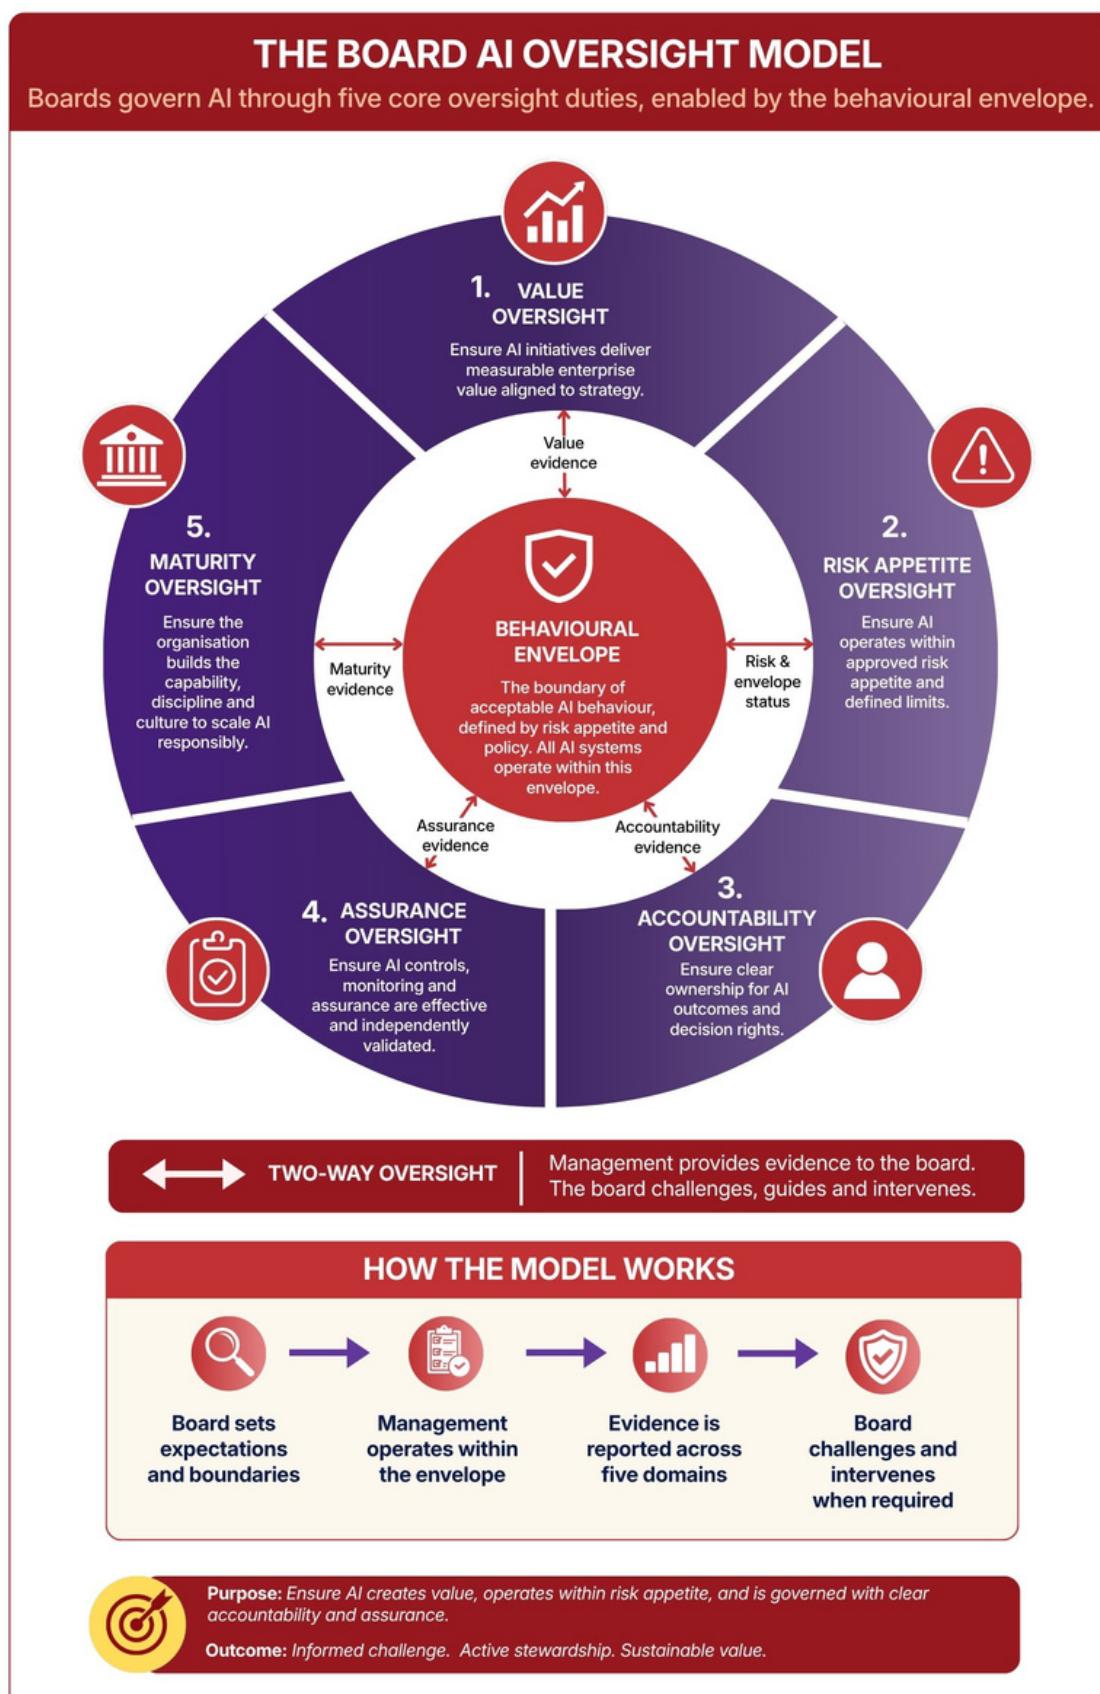


Figure 3: Board AI Regulation Question Tree

Figure 4 sets out the five oversight duties that provide the structural framework through which boards should organise their regulatory oversight inquiry.

**Figure 4: AI Board Oversight Model**

Regulatory Exposure

Where are our highest regulatory exposure points across AI use cases, and how are these identified and prioritised?

Risk Classification and Boundaries

Which of our AI systems would be considered high-risk under current regulatory frameworks, and what additional controls are in place for these systems?

Evidence of Control

What evidence do we have that controls are operating effectively in practice, not just defined in policy? Can that evidence be produced quickly if a regulator requests it?

Monitoring and Incident Response

How are we monitoring AI systems over time, and how are incidents, drift or unexpected behaviour identified, escalated and managed?

Accountability

Who is accountable for regulatory compliance of each AI system, and how is that accountability demonstrated rather than assumed?

Regulatory Change Readiness

How are we tracking regulatory developments, and how quickly can the operating model adapt in response to material changes in regulatory expectations?

What Leaders Should Do Now

Regulatory readiness requires building capability over time. The following sequence provides a practical starting point. The order matters: each step creates the foundation for the next.

Map regulatory exposure.

Identify which AI systems are in use, what regulatory domains they intersect with, and where exposure is highest. Without this view, governance cannot be targeted.

Define governance and accountability.

Establish clear ownership for AI systems and regulatory obligations. Good governance means named individuals are accountable for specific outcomes, not shared responsibility across teams.

Translate requirements into the operating model.

Regulatory obligations must become decision rights, roles, policies and processes embedded in how the organisation operates. Documentation that sits outside day-to-day operations does not constitute compliance.

Implement controls.

Design and embed controls across the AI lifecycle. Controls must be proportionate to the risk level of each system and must operate in production, not only at the point of approval.

Establish evidence and monitoring.

Build the capability to generate, retain and present evidence that controls are operating effectively. Monitoring must be continuous, not periodic, for systems with material regulatory exposure.

Enable assurance.

Independent validation from internal audit and, where appropriate, external review provides the board with confidence that the operating model is functioning and that regulatory obligations are being met.

This is not a one-time exercise. Regulatory expectations will continue to evolve, and the operating model must adapt accordingly.

Failure Modes

Common failure modes are already emerging across organisations attempting to respond to AI regulation. Each reflects a predictable pattern of governance design that looks complete but fails under scrutiny.

Policy-Led Compliance

Organisations rely on policies without embedding them in operational processes, creating a gap between intent and reality. In practice, this appears as well-documented frameworks that cannot be evidenced during regulatory scrutiny. When challenged, teams point to the policy document rather than demonstrating that the policy is operating.

Fragmented Accountability

Responsibility is distributed without clear ownership, leading to inconsistent decisions and weak governance. This often results in multiple teams assuming others are accountable, with no single point of responsibility when something goes wrong. Regulators will ask who is responsible. The answer must be immediate and specific.

Framework Adoption Without Translation

Frameworks are adopted without being integrated into the operating model, resulting in theoretical alignment but practical gaps. Organisations appear compliant on paper but lack execution capability. The framework describes what good looks like. Without translation into roles, processes and controls, it does not govern behaviour.

Absence of Evidence

Controls exist but cannot be demonstrated, leaving organisations exposed under regulatory scrutiny. When challenged, teams are unable to produce consistent monitoring data or proof of control effectiveness. Evidence is the currency of regulatory confidence. If it cannot be produced, the control is not functioning in any meaningful governance sense.

Static Governance Models

Governance structures fail to adapt to regulatory change, creating ongoing exposure. Over time, controls become outdated while regulatory expectations continue to evolve. An organisation that built its governance model in response to current regulation, and then stopped, will fall behind as the regulatory landscape continues to develop.



Closing Insight

AI regulation is not a legal problem to solve. It is a governance capability to build.

Regulation defines expectations. Organisations must build the capability to meet those expectations consistently and at scale.

The governance model required to achieve this is not new. It is the same model that underpins effective AI transformation more broadly. It requires clear governance structures, disciplined portfolio management, defined accountability, embedded controls, evidence and monitoring, and independent assurance.

These are the capabilities described across this series. Regulatory readiness is not a separate effort. It is the outcome of building these capabilities effectively.

Boards do not need to interpret regulation. They need to ensure that the organisation can operate within it, adapt to change and demonstrate that this is happening.

The organisations that succeed will not be those that respond to regulation reactively. They will be those that build governance, operating model and assurance as an integrated and adaptive capability, one that creates value, controls risk and meets regulatory expectations not as a compliance exercise, but as a natural product of how they govern AI.

References

European Union. Artificial Intelligence Act (EU AI Act). Official Journal of the European Union, 2024.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST, 2023.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative AI Profile. NIST, 2024.

International Organization for Standardization. ISO/IEC 42001: Artificial Intelligence Management System. ISO, 2023.

Organisation for Economic Co-operation and Development. OECD AI Principles. OECD, updated 2024.

Australian Government Department of Industry, Science and Resources. Voluntary AI Safety Standard. Australian Government, 2024.

Australian Government Department of Industry, Science and Resources. Safe and Responsible AI in Australia. Australian Government, 2024.

Australian Prudential Regulation Authority. Prudential Standard CPS 230: Operational Risk Management. APRA, 2023.

Australian Securities and Investments Commission. Corporate Governance and Artificial Intelligence: Director and Officer Responsibilities. ASIC, 2024.

Deloitte. State of Generative AI in the Enterprise. Deloitte AI Institute, 2025.

EY. Responsible AI Pulse Survey. EY Global, 2025.

KPMG. Trustworthy AI: A Risk-Based Approach to AI Governance. KPMG, 2025.

PwC. Responsible AI Survey. PwC, 2025.

World Economic Forum. Responsible AI Playbook. WEF, 2025.

Stanford Institute for Human-Centered Artificial Intelligence. AI Index Report. Stanford University, 2025.

Casepoint. AI Regulatory Readiness: Enterprise Preparedness Survey. Casepoint, March 2026.

Australian Prudential Regulation Authority. APRA Communication on Artificial Intelligence Governance and Operational Risk. APRA, April 2026.

Office of the Australian Information Commissioner. Guidance on Privacy and Artificial Intelligence. OAIC, 2024.

Council of Europe. Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. Council of Europe, 2024.

State of Colorado. Colorado Artificial Intelligence Act (SB 24-205). State of Colorado, 2024.

Disclaimer: *This paper does not constitute legal advice. The regulatory landscape described is subject to ongoing change. Specific regulatory obligations, legal exposure and compliance requirements should be assessed with qualified legal and compliance advisers.*

About the Author

Manoj Tavarajoo has spent over two decades working with boards, executives and senior leaders across enterprise, digital and AI transformation. His work sits at the intersection of strategy, operating model design, governance and execution, where transformation ambition becomes measurable enterprise capability.

He is the author of *Leading the AI Transformation*, *The AI Operating Model Playbook* and *The AI Governance Imperative*, together forming a practical leadership trilogy on how organisations lead AI, organise for AI and govern AI at enterprise scale.

Manoj works through MyConsultancy, an independent advisory and transformation practice based in Australia, and is open to senior embedded leadership opportunities partnering with boards and executive teams navigating AI strategy, operating model design and governance at enterprise scale.




[@manojtavarajoo](https://www.linkedin.com/company/manojtavarajoo)

About MyConsultancy

MyConsultancy works with boards and executive teams navigating the distance between AI ambition and operating reality. The practice specialises in AI strategy, operating model design, governance and enterprise transformation, helping organisations build the structures, decision-making disciplines and execution capability needed to scale AI responsibly.



www.myconsultancy.com.au



MyConsultancy