Evolve.i

EVOLVE.INSIGHTS

Responsible AI by Design: What Every Organisational Role Needs to Do

A practical, role-by-role governance framework that helps Australian Boards, CEOs, executives, leaders and staff introduce, manage and monitor AI responsibly, without slowing productivity, efficiency or innovation.

A simple governance principle

Use the lightest control that is appropriate to the risk, and strengthen oversight as the potential impact on people, services, finances and data increases.

How to use this framework

Responsible AI is not a technology project. It is an operating discipline. This framework provides a proportionate operating model for governing AI across an organisation. It is designed for Boards, executives, business leaders and operational teams that want to capture the value of AI while maintaining clear accountability, human oversight and evidence of responsible practice.

> **Regulatory position: July 2026**
>
> The NSW AI Assessment Framework is mandatory for NSW Government agencies. Other organisations can use it as a strong governance benchmark, but should not describe it as a universal legal requirement. Across Australia, existing laws already apply to AI, including privacy, consumer, workplace, anti-discrimination, corporations, intellectual property, safety and cyber security obligations. The Australian Government has also issued six essential AI practices. From 10 December 2026, relevant APP entities will have additional privacy-policy transparency obligations for certain automated decisions using personal information.

The governance model in one page

Responsible AI does not require a new bureaucracy for every tool. It requires one repeatable governance loop, clear decision rights and monitoring that is proportionate to risk.

1. REGISTER Record every approved AI system and material use case.	**2. ASSESS** Classify the use by impact, data sensitivity and decision influence.	**3. APPROVE** Apply the right approval pathway and document accepted risk.
4. USE Train users, preserve human control and follow approved conditions.	**5. MONITOR** Measure performance, benefits, incidents, complaints and drift.	**6. IMPROVE** Adjust, pause, retire or scale the use based on evidence.

The minimum governance set

- **AI policy and staff guidance:** Define permitted, restricted and prohibited uses; data-handling rules; disclosure expectations; and escalation pathways.
- **AI register:** Record the tool, use case, owner, vendor, data, risk tier, approvals, review date, performance and incidents.
- **Risk and impact assessment:** Assess purpose, affected stakeholders, decision influence, privacy, cyber security, fairness, explainability, legal obligations and business continuity.
- **Human oversight:** Name the person accountable for reviewing outputs, intervening, overriding or stopping the system.

- **Vendor assurance:** Confirm data use, retention, model training, hosting, subcontractors, security, audit rights, incident notification and exit arrangements.
- **Monitoring and reporting:** Track value and risk together, with review frequency determined by the risk tier.

1. Classify AI use by risk, not by brand or tool

The same AI product can be low risk in one context and high risk in another. Classify the use case by what the system does, the information it handles, who is affected and the consequence if it is wrong.

Tier	Typical use	Risk indicators	Minimum approval	Review
1 LOW	Drafting, summarising or ideation using public, approved or non-sensitive information.	No material decision; no personal or confidential data; output is reviewed before use.	Manager or use-case owner; register entry.	6 to 12 months
2 MODERATE	Internal analysis, workflow support or recommendations that influence work but do not determine high-impact outcomes.	Internal or confidential data; operational reliance; moderate financial or service impact.	Business owner plus technology, data and privacy review.	Quarterly
3 HIGH	AI affecting employment, access to services, financial outcomes, safety, health, eligibility or customer rights.	Personal or sensitive data; vulnerable people; significant decision influence; external-facing or autonomous action.	Executive AI Committee; formal impact, privacy and cyber assessments; documented human oversight.	Monthly or continuous
4 CRITICAL	Use that cannot be made lawful, safe, fair, contestable or adequately controlled.	Fully automated high-impact decisions without meaningful review; covert or unlawful use; unacceptable safety, discrimination or privacy risk.	Do not deploy, or escalate to Board and CEO with specialist legal and assurance advice.	Continuous or immediate

Fast escalation triggers

Escalate an AI use case whenever it uses personal or sensitive information, influences a decision about a person, operates without timely human review, connects to a high-risk system, creates external content or advice, or could materially affect safety, legal rights, funding, employment, reputation or service access.

Pre-use questions for every AI system

- What business or community outcome is the AI intended to improve, and is AI the most appropriate solution?
- Who could be affected, including employees, customers, clients, participants, vulnerable people or communities?
- What information will enter, train, configure or be generated by the system?
- Does the AI make, recommend, prioritise or substantially influence a decision?
- What must a human review, approve, override or explain?
- How will accuracy, bias, reliability, security and performance be tested?
- What will be disclosed to users and affected people, and how can they ask questions or challenge an outcome?
- What evidence will demonstrate that benefits continue to outweigh risks?

2. Role-by-role accountability

Accountability should follow existing organisational authority. The Board governs; the CEO is accountable for the organisational system; executives own their risk domains; business owners are responsible for each use case; and every user remains accountable for how they use AI outputs.

Role	Before introduction	Ongoing monitoring	Cadence
Board / Board Risk & Audit Committee	Confirm AI risk appetite, ethical principles, prohibited uses and the level of assurance expected for high-impact AI.	Receive a concise AI dashboard; challenge whether benefits outweigh risk; review material incidents, high-risk deployments, residual risk and regulatory readiness.	Quarterly; immediately for critical incidents.
Chief Executive Officer	Appoint the executive accountable for AI governance; approve the governance model, resources and culture expectations; ensure AI aligns with strategy and organisational values.	Report to the Board; review the enterprise AI portfolio, high-risk approvals, material exceptions, adoption, value, workforce impact and incidents.	Monthly executive review; quarterly Board reporting.
AI Governance Committee / Executive AI Sponsor	Operate the governance process; classify risk; coordinate legal, risk, privacy, technology, people and business input; approve or escalate use cases.	Maintain the AI register, decision records, risk treatments, review calendar, exceptions, incidents and improvement actions.	Monthly; more often during implementation.

Role	Before introduction	Ongoing monitoring	Cadence
Chief Operating Officer / Transformation Lead	Ensure AI is introduced through redesigned processes, clear controls, service continuity and measurable operating outcomes.	Track process performance, quality, cycle time, rework, adoption, customer impact, dependencies and fallback arrangements.	Monthly.
Chief Financial Officer	Validate the business case, funding, procurement controls, financial delegation, value assumptions and financial and reporting risks.	Monitor total cost, realised savings, productivity, duplicated licences, financial errors, fraud exposure and return on investment.	Monthly or quarterly according to materiality.
Chief People Officer / HR	Assess workforce impact, consultation, role redesign, capability needs, acceptable use, employee monitoring, fairness and employment decision risks.	Monitor training completion, adoption, workload, role changes, grievances, inclusion, psychological safety and AI use in recruitment or performance decisions.	Monthly during change; quarterly thereafter.
CIO / CTO / CISO / Data Lead	Assess architecture, integration, identity, access, data quality, privacy controls, cyber security, vendor technical assurance, logging and resilience.	Monitor security events, data leakage, access, system performance, model changes, integrations, technical incidents, vendor updates and business continuity.	Continuous monitoring; monthly reporting.
Legal, Risk, Compliance & Privacy	Identify applicable laws, contractual requirements and sector obligations; determine impact assessments, notices, consent, record keeping, contestability and reporting obligations.	Track regulatory changes, privacy and legal exceptions, complaints, breaches, residual risk, control effectiveness and remediation.	Quarterly; immediate escalation for material issues.
Procurement / Vendor Management	Apply AI-specific due diligence and contract clauses before purchase, renewal or material change.	Monitor vendor performance, data practices, subcontractors, certifications, incidents, model or terms changes, audit evidence and exit readiness.	At procurement, renewal and material change.
Business / Product / Process Owner	Define the intended use, affected stakeholders, success measures, acceptance criteria, human oversight, approved users and operating procedure.	Review outputs, accuracy, bias, complaints, overrides, benefits, failure modes, user feedback and whether the use remains within scope.	Monthly for high risk; quarterly for moderate risk.
Leadership / People Managers	Translate policy into team practice; approve appropriate low-risk use; ensure training, supervision and escalation;	Review team usage, quality, exceptions, incidents, capability gaps and improvement opportunities.	Monthly team check-in.

Role	Before introduction	Ongoing monitoring	Cadence
	prevent unapproved or shadow AI use.		
Internal Audit / Independent Assurance	Provide independent review of governance design, risk classification, control operation and evidence for material AI use.	Test a risk-based sample of systems, approvals, data controls, monitoring, human intervention, vendor assurance and remediation closure.	Annual plan plus targeted reviews.
All Staff, Contractors & Volunteers	Use only approved tools and approved data; verify outputs; protect confidential information; disclose AI use where required; never delegate personal accountability to AI.	Report errors, harmful or biased outputs, privacy or security concerns, unexpected behaviour and unapproved AI use promptly.	At every use; annual refresher training.

Accountability should be confirmed for each use case through a practical RACI (Accountable, Responsible, Consulted). The exact allocation will vary by organisation size and risk.

3. Monitor value and risk together

Good governance should not measure compliance in isolation. Each AI system should have a small set of indicators that show whether it remains useful, controlled and aligned with its approved purpose.

Monitoring domain	Example measures	Owner	Escalation signal
Business value	Time saved; cost avoided; cycle time; service capacity; revenue or funding impact; quality improvement; user adoption.	Business owner / CFO / COO	Benefits materially below business case; increased rework; dependency without resilience.
Accuracy & quality	Error rate; hallucination rate; acceptance rate; override rate; rework; false positives and negatives; output consistency.	Business owner / Technology	Performance below acceptance threshold or unexplained change over time.
Fairness & people impact	Outcome differences between groups; complaints; accessibility; workforce impact; adverse decisions; stakeholder feedback.	CPO / Risk / Business owner	Evidence of disadvantage, discrimination, unsafe impact or unresolved stakeholder concern.
Privacy, data & cyber	Personal and sensitive data use; leakage; access anomalies; retention; data quality; security events; unauthorised tools.	Privacy / CISO / Data lead	Data breach, unauthorised disclosure, material access failure or unapproved data use.
Human control	Review completion; override capability; intervention time;	Use-case owner / Risk	Human review is bypassed, ineffective,

Monitoring domain	Example measures	Owner	Escalation signal
	decision traceability; appeal and redress performance.		delayed or unable to explain or alter an outcome.
Vendor & technical change	Model and version change; terms or data-use change; uptime; incidents; subcontractors; control attestations.	Technology / Procurement	Material vendor change without reassessment or failure to provide required evidence.
Compliance & assurance	Assessments current; approvals complete; training completion; audit findings; actions overdue; privacy notices current.	AI Committee / Compliance	Expired approval, overdue high-risk review, repeated control failure or material audit finding.

Suggested Board and executive dashboard

- Number of registered AI systems and use cases by risk tier.
- New, changed, paused and retired AI uses during the period.
- High-risk systems with current impact assessments, approvals and named human owners.
- Material incidents, complaints, privacy or cyber events, and overdue corrective actions.
- Performance against benefit, accuracy, fairness, quality and service measures.
- Workforce adoption, training completion, unapproved AI use and capability gaps.
- Vendor and model changes and upcoming regulatory or assurance actions.
- Decisions required from the Board or executive team.

Keep reporting decision-focused

The Board does not need prompt-level detail or a catalogue of every experiment. It needs assurance that material uses are visible, risk is owned, human control is effective, benefits are being realised and issues are escalated early.

4. A practical 90-day implementation plan

Most organisations can establish a credible governance foundation in 90 days by building around existing risk, technology, privacy, people and procurement processes rather than creating a parallel compliance system.

DAYS 1 TO 30 | Establish visibility and boundaries

- Nominate the executive AI sponsor and establish a small cross-functional AI Governance Committee.
- Approve interim staff guidance: approved tools, prohibited data, output verification, disclosure and incident reporting.
- Create the AI register and complete an initial discovery of tools already in use, including AI embedded in existing software.
- Define risk tiers, escalation triggers and an interim approval pathway.
- Identify existing high-impact or personal-information uses requiring immediate review.

DAYS 31 TO 60 | Embed proportionate controls

- Approve the AI policy, role accountabilities and risk appetite.
- Introduce a short use-case intake and assessment integrated with procurement, privacy, cyber and project governance.
- Complete vendor assurance for priority systems and establish AI-specific contract clauses.
- Define acceptance tests, monitoring measures, human oversight and fallback procedures for moderate and high-risk uses.
- Train executives, leaders, technology teams and general users according to their responsibilities.

DAYS 61 TO 90 | Demonstrate evidence and improve

- Review the first portfolio dashboard with the executive team and Board committee.
- Complete targeted assurance of high-risk use cases and close urgent gaps.
- Update privacy notices and prepare for the 10 December 2026 automated decision transparency requirements where applicable.
- Establish ongoing review dates, incident simulations and an annual independent assurance plan.
- Use evidence to scale valuable low-risk uses and pause or redesign uses where controls or benefits are weak.

Scale the model to organisation size

Organisation	Governance arrangement	Minimum review
Small	CEO or owner as accountable executive; a nominated business owner; external legal, privacy and technology advice when needed. One register and one quarterly review.	Quarterly portfolio review; annual policy review.
Medium	Cross-functional AI Governance Committee led by an executive sponsor; defined product and process owners; internal risk, people, finance and technology review.	Monthly committee; quarterly executive and Board dashboard.
Large / high-regulation	Formal AI governance function or committee; model risk, privacy, cyber, legal, people and independent assurance integrated with enterprise risk and audit.	Continuous high-risk monitoring; monthly committee; quarterly Board oversight.

5. Practical templates

AI use-case intake: minimum information

Complete the following for each material use case:

- Use-case name and owner
- Business purpose and intended benefit
- AI tool, vendor and model
- Users and affected stakeholders
- Information used, including personal and sensitive data
- How the output influences decisions or actions
- Human review, intervention and escalation
- Risk tier and required assessments
- Acceptance criteria and testing evidence
- Monitoring measures and review date
- Approval decision, conditions and residual risk owner

> **AI register: recommended fields**
>
> System or use case; owner; business purpose; vendor and model; risk tier; affected stakeholders; data classification; personal and sensitive information; decision influence; human oversight; assessments completed; approval and conditions; acceptance criteria; monitoring metrics; incidents and complaints; vendor and model changes; last review; next review; status; retirement date.

Stop and pause criteria

- Performance falls below an approved safety, accuracy, quality or fairness threshold.
- The system acts outside its approved purpose, users or data boundaries.
- A material privacy, cyber, safety, legal, discrimination or integrity incident occurs.
- Meaningful human review or intervention is unavailable or ineffective.
- The vendor, model, data source, terms or integration changes materially without reassessment.
- Complaints or affected-stakeholder feedback indicate a systemic issue.
- The system no longer produces sufficient benefit to justify its risk, cost or complexity.

Board paper: recommended headings

- Decision sought
- Purpose and strategic alignment
- Use case and affected stakeholders
- Benefits and success measures
- Risk tier and assurance completed
- Human oversight and contestability
- Privacy, cyber, data and vendor controls
- Residual risks and accountable owner
- Monitoring and reporting plan
- Recommendation and conditions of approval

6. Alignment with NSW and Australian guidance

This framework translates the common expectations in NSW and Australian guidance into a practical organisational operating model. It does not replace sector-specific legal advice, contractual obligations, regulator guidance or the NSW AI Assessment Framework process for NSW Government agencies.

Government expectation	How this framework responds	Primary evidence
NSW: AI is appropriate and creates community and customer benefit.	Business purpose, alternatives, affected stakeholders, benefit measures and proportional risk tiering.	Use-case assessment; business case; benefits dashboard.
NSW: Fairness and bias are considered.	Stakeholder impact, fairness measures, human review, complaint and redress pathways.	Impact assessment; testing; monitoring; complaint records.

Government expectation	How this framework responds	Primary evidence
NSW: Privacy and security are protected.	Data classification, PIA trigger, cyber review, vendor due diligence, access, retention and leakage monitoring.	PIA; security review; contract; logs; incident records.
NSW: Transparency and review mechanisms exist.	Clear notices, AI register, decision traceability, disclosure and accessible challenge pathways.	Privacy notice; system information; appeals and feedback process.
NSW: Accountability and human intervention are clear.	Named accountable executive, use-case owner, RACI, approval record and stop and override authority.	Governance charter; approval; operating procedure; training.
Australian Government: six essential AI practices.	Accountability; impact planning; AI-specific risk management; information sharing; testing and monitoring; human control.	Policy; register; assessments; testing; dashboard; intervention plan.
Privacy Act: automated decision transparency from 10 December 2026 for relevant APP entities.	Identify automated decision uses that significantly affect rights or interests and prepare the required privacy-policy disclosures.	ADM inventory; legal assessment; updated privacy policy.

Key source material

- **NSW AI Assessment Framework:** open official source
- **NSW Artificial Intelligence Ethics Policy:** open official source
- **Digital NSW: Understanding Responsibilities in AI Practices:** open official source
- **Australian Government: Guidance for AI Adoption:** open official source
- **OAIC: Privacy and commercially available AI products:** open official source
- **OAIC: Automated decision-making transparency consultation:** open official source

Important: This framework is general information and is not legal, privacy, cyber security, employment, financial or sector-specific advice. Organisations should obtain appropriate specialist advice for their circumstances and monitor regulatory change.

Governance should enable AI to evolve, not prevent it.

The goal is not to remove every risk. It is to make risk visible, assign ownership, maintain human control and use evidence to improve both performance and trust.

How Evolve.i can support

AI readiness and governance diagnostics • Board and executive advisory • AI policy and controls • Operating model and process redesign • Staff guidance and capability • Implementation and assurance

connect@evolvei.com.au • www.evolvei.com.au