

A Simple AI Governance Readiness Checklist

1. Leadership & Accountability

- 1.1 Has a C-suite sponsor formally accepted ultimate accountability for AI governance (charter, budget, KPIs)?
- 1.2 Are roles, responsibilities, and decision rights for AI risks documented and visible across the organization?
- 1.3 Does an AI Ethics/Governance board meet on a scheduled cadence with cross-functional voting members and published minutes?

2. Responsible-AI Policy & Principles

- 2.1 Are written RAI principles covering fairness, accountability, transparency, privacy, security, robustness, and sustainability approved and published?
- 2.2 Is there a controlled process to translate principles into enforceable policy (incl. change control and version history)?
- 2.3 Have policies been mapped to ISO 42001 clauses, NIST AI RMF functions, and relevant regional regulations?

3. Risk Management & Regulatory Compliance

- 3.1 Has the organization completed an AI risk register aligned to NIST AI RMF (Map-Measure-Manage-Govern) with risk owners assigned?
- 3.2 Are documented controls in place for all “high-risk” uses under the EU AI Act (or equivalent local laws)?
- 3.3 Are AI Risk Assessments completed for all new AI systems during design or procurement?
- 3.4 Are AI System Impact Assessments completed for all new AI systems during design or procurement?
- 3.5 Is there an independent internal or external audit mechanism that reviews AI systems at least annually?

4. Data & Model Governance

- 4.1 Data provenance & consent: Is the origin, licensing, and consent status of all training data tracked and reviewable?
- 4.2 Bias & quality testing: Are datasets and models routinely tested for representational bias, drift, and data quality issues?
- 4.3 Versioning & lineage: Do you maintain immutable version control for datasets, models, and experiments?
- 4.4 Access controls: Are least-privilege and segregation-of-duties enforced for data scientists vs. production teams?

5. Technical Robustness & Security

- 5.1 Have threat-modelling and adversarial-robustness assessments been performed for each model before deployment?
 - 5.2 Are secure coding, Software Bill of Material, and vulnerability-scanning practices mandated for AI code and pipelines?
 - 5.3 Is runtime monitoring in place to detect and respond to model degradation, prompt injection, or data poisoning attacks?
- 6. Stakeholder Engagement & Transparency**
- 6.1 Has a stakeholder analysis identified internal & external users, impacted communities, and regulators for every AI use case?
 - 6.2 Do you publish plain-language AI system summaries describing purpose, limitations, and contact points?
 - 6.3 Is there a two-way channel (portal, hotline) for inquiries, appeals, and feedback on AI system outcomes?
- 7. Organizational Culture & Capability**
- 7.1 Has an organizational readiness survey measured trust, understanding, and ethical climate regarding AI?
 - 7.2 Do mandatory RAI trainings exist for all employees, with role-based deep dives for designers, developers, and auditors?
 - 7.3 Are communities of practice or centers of excellence actively sharing lessons and updating playbooks?
- 8. AI Lifecycle Integration & Third-Party Management**
- 8.1 Are ethical “gates” built into SDLC phases (design, data prep, modeling, validation, deployment, retirement) with sign-off checklists?
 - 8.2 Do procurement & vendor-management processes require suppliers to attest to compliance (e.g., ISO 42001, SOC 2, OWASP LLM Top 10)?
 - 8.3 Are automated CI/CD pipelines configured to block promotion when governance tests fail?
- 9. Incident Response & Redress**
- 9.1 Is there an AI-specific incident-response plan covering detection, triage, containment, root-cause, and public communication?
 - 9.2 Have clear redress / appeal mechanisms been defined for individuals adversely affected by AI decisions?
 - 9.3 Are post-incident reviews feeding lessons learned back into policies, datasets, and models within a fixed SLA?
- 10. Metrics, Monitoring & Continuous Improvement**
- 10.1 Do defined KPIs track adherence to each RAI principle (e.g., disparity impact ratio, explanation coverage, mean-time-to-mitigate)?

- 10.2 Are dashboards and periodic reports delivered to the board, showing trend lines and control effectiveness?
- 10.3 Does a formal review cycle (at least yearly) update policies, training, and technical controls based on monitoring data and new regulations?
- 10.4 Is a budget and roadmap allocated for ongoing capability maturation (tooling upgrades, staff development, external benchmarking)?