AI GOVERNANCE READINESS CHECKLIST

Is your organisation ready to deploy AI responsibly?

15 questions your legal, risk, and engineering teams need to align on before you go live. Score each item: 2 = fully in place · 1 = partial / in progress · 0 = not yet addressed.

01 · GOVERNANCE & ACCOUNTABILITY	0	1	2
Ownership is assigned A named individual or committee owns AI risk, with a mandate that includes budget and veto authority.	0	1	2
Policies exist and are enforced Written AI use policies cover acceptable use, prohibited use, and escalation paths — not just principles.	0	1	2
Model inventory is maintained Every AI system in production is logged: purpose, owner, training data provenance, and last review date.	0	1	2
Third-party AI is governed Vendor AI (SaaS, APIs, embedded models) is subject to the same risk controls as internally built systems.	0	1	2

02 · DATA & PRIVACY	0	1	2
Personal data flows are mapped You know which AI systems process personal information and under which legal basis (PIPEDA / FIPPA / GDPR).	0	1	2
Training data is documented Data sources, consent status, and any third-party licensing constraints are recorded for each model.	0	1	2
Data retention and deletion apply to AI outputs Outputs, logs, and embeddings are included in your retention schedule — not treated as non-personal.	0	1	2

03 · RISK & COMPLIANCE	0	1	2
A risk classification exists AI use cases are classified by risk tier. High-risk applications trigger additional controls before deployment.	0	1	2
Regulatory mapping is current Relevant frameworks (ISO 42001, EU AI Act, NIST AI RMF, OSFI B-10) are mapped to your actual controls.	0	1	2
Bias and fairness are tested High-stakes models (hiring, lending, triage, access decisions) are evaluated for disparate impact before launch.	0	1	2
An incident response plan covers AI Model failure, hallucination in production, and data leakage through AI outputs are covered in your IR playbook.	0	1	2

04 · ARCHITECTURE & OPERATIONS	0	1	2

Human oversight is designed in For consequential decisions, a human review step is architecturally enforced — not just recommended in policy.	0	1	2
Models are monitored in production Drift, performance degradation, and anomalous outputs trigger alerts with defined SLAs for investigation.	0	1	2
Retrieval-augmented systems are auditable For RAG deployments, retrieved sources are logged and traceable to each generated output.	0	1	2
An exit and rollback plan exists You can disable or roll back any AI system within a defined timeframe without breaking dependent processes.	0	1	2

SCORING GUIDE	24–30	15–23	8–14	0–7
Add up your scores. Each item: 2 = fully in place · 1 = partial · 0 = not yet.	Strong foundation. Focus on gaps.	Developing. Prioritise controls.	Significant exposure. Architectural review needed.	High risk. Governance programme required.

Scored below 20? Not sure where to start?

Manseur Advisory provides independent AI governance and architecture counsel — no firm overhead, senior practitioners only. Ready to go deeper? manseur.com · Book a conversation