



# AI Governance Readiness Checklist

## Human oversight you can prove — not just promise.

Regulatory pressure on AI is moving from guidance to enforcement. The EU AI Act begins imposing real penalties on high-risk AI systems from August 2026. ISO/IEC 42001, the international AI management standard, is increasingly showing up as a procurement requirement. Most organizations can't yet answer a simple question with evidence: when an AI agent made a recommendation, who reviewed it, and what did they decide?

*How to use this: work through the five sections below with your team. Check off what's genuinely true today, not what's planned. Unchecked items are your starting point, not a failing grade.*

### G Governance

*Decision rights, policies, and oversight structures that make clear who can approve what an AI agent recommends — defined before deployment, not reconstructed after an incident.*

- ☐ Decision rights are documented: it's clear who can approve, override, or reject an AI-generated recommendation.
- ☐ Oversight responsibilities were defined before deployment, not created after an incident.
- ☐ There's a named owner accountable for each AI system's outputs.
- ☐ Escalation paths exist for when an AI recommendation is contested or unclear.

### E Ethics

*Fairness, transparency, and avoiding harm built into how recommendations are reviewed and acted on — not a values statement bolted on after the system ships.*

- ☐ Recommendations are reviewed for fairness and bias before being acted on.
- ☐ The basis for an AI recommendation is explainable to the people who rely on it.
- ☐ Ethical guardrails are built into the review workflow, not just written in a values statement.
- ☐ Potential harms were considered during design, not only after deployment.

### R Responsible AI

*AI applied where it genuinely adds value, with explicit limits on what it shouldn't decide alone. Always advisory over consequential decisions — never autonomous.*

- ☐ There's an explicit, written list of decisions the AI is not permitted to make alone.
- ☐ AI is used where it adds genuine value — not adopted for its own sake.
- ☐ Consequential decisions always have a human reviewer before action is taken.

- ☐ The system's limits are known and communicated to the people using it.

## D Data Quality

*The CPMAI-aligned discipline of understanding, auditing, and validating the data an AI system learns from — before anyone trusts what it recommends.*

- ☐ The data an AI system learns from has been audited for completeness and accuracy.
- ☐ Known biases or gaps in the training/input data are documented.
- ☐ Data provenance is tracked — you know where the data came from.
- ☐ Data quality is revisited periodically, not validated once and forgotten.

## A Audits

*A timestamped, tamper-evident record of every AI recommendation and every human decision — ready to show a regulator, a board, or a client, not assembled after the fact.*

- ☐ Every AI recommendation is logged with a timestamp.
- ☐ Every human decision made in response to an AI recommendation is also logged.
- ☐ Audit records are tamper-evident, not just a spreadsheet anyone can edit.
- ☐ You could produce this record on short notice for a regulator, board, or client today.

---

***AI provides the analysis. The person owns the decision.***

### **Assess your AI governance readiness**

A 30-minute conversation to understand where your PMO or organization stands, and what audit-readiness would actually take.

**Get in touch: [info@nextmtech.com](mailto:info@nextmtech.com) | Book a 30-min call: [calendly.com/d-haddad-nextmtech/30min](https://calendly.com/d-haddad-nextmtech/30min)**