AI Governance Readiness Checklist

ISO/IEC 42001 AI management system requirements and EU AI Act classification, risk, impact and oversight obligations

ISO/IEC 42001 · EU AI Act · AI Governance · AI Risk Management

Two laws sit behind AI governance. ISO/IEC 42001 certifies your organisation's AI management system as a whole. The EU AI Act regulates each AI system individually. This checklist covers what both actually require, so a readiness gap is something you find on your own terms.

A EU AI Act Classification

- [ ] Every AI system in use has been checked against banned-use categories first
- [ ] Each system's risk category (high-risk, transparency-only, or out of scope) is determined and recorded with legal reasoning
- [ ] Your role for each system is documented — provider/manufacturer or deployer — since the law imposes different duties on each
- [ ] A system that has been substantially modified, rebranded, or repurposed has been re-checked, not assumed still correctly classified

B ISO/IEC 42001 Management System

- [ ] An AI policy exists, is version-controlled, and has been published to affected staff
- [ ] Named, accountable owners exist for AI governance, risk, compliance, data and security
- [ ] Measurable AI governance objectives are set and tracked, not just stated
- [ ] AI literacy and competence records exist for relevant roles
- [ ] Internal audits of the AI management system run to a programme, and management reviews are held on a schedule

C Compliance & Applicability

- [ ] Every applicable requirement, per framework, has a recorded position: met, partly met, not met, or not yet assessed
- [ ] Every excluded requirement has a documented reason — an unexplained exclusion is an easy audit failure
- [ ] Evidence is attached per requirement, not held informally or only in someone's inbox
- [ ] A Statement of Applicability exists for ISO/IEC 42001, listing every control applied or excluded with reasons

D Impact Assessment

- ☐ Affected people are identified for each system, including anyone vulnerable (children, patients, applicants)
- ☐ Harm areas are assessed: privacy, bias, discrimination, safety, financial loss, dignity, accessibility, employment
- ☐ Each harm is rated twice — before safeguards, and again with safeguards counted — not just once
- ☐ Where residual risk is unacceptable, treatment actions exist with an owner and a deadline

E Controls, Evidence & Human Oversight

- ☐ Controls are mapped to every requirement they satisfy across every framework, not documented once per framework
- ☐ Each control is assessed three ways: well designed, actually in place, and working — not assumed from one check
- ☐ Evidence carries a validity period, so stale proof is visible rather than quietly assumed current
- ☐ For high-risk systems, human oversight is recorded as real events — reviews, approvals, overrides — not asserted as a policy statement

F AI Risk, Incidents & Change

- ☐ AI-specific risks (biased output, data leakage, model drift, over-reliance, vendor dependency) are scored in the same risk register and framework as every other business risk
- ☐ An incident process exists that starts the regulatory notification clock the moment an AI incident is logged, not once someone remembers the deadline
- ☐ A change process catches model, provider, purpose, data-access or autonomy changes and triggers the right re-review — risk, impact, or full reclassification
- ☐ A dashboard or system profile shows, per AI system, its classification, risk, controls, oversight and open actions in one place