AI Governance Policy Starter Pack

Three starter policies you can adapt: acceptable use, procurement, and human oversight. Each is a skeleton with the decisions you need to make filled in as prompts - keep it short enough that people actually read it.

1. AI Acceptable Use Policy

Purpose and scope

States who the policy covers (all staff, contractors) and which tools it covers: approved AI tools, AI features inside other software, and personal accounts used for work. PROMPT: name your approved-tool list and where it lives.

Permitted use

Drafts, summarization, code assistance, research - with the rule that a named human reviews and takes responsibility for anything AI-assisted that leaves the organization.

Prohibited use

PROMPT: your specific red lines. Common ones: personal data of clients or staff into unapproved tools; automated decisions about individuals without human review; representing AI output as verified fact; customer-facing content without review.

Data handling

Which data classes may never be entered into which tools, and what to do instead (approved deployments, anonymization, or vendor DPAs).

Disclosure

When staff must tell colleagues, clients, or customers that content was AI-assisted.

Incidents

Who to tell when confidential data is entered into an unapproved tool, and by when.

2. AI Procurement Policy

Threshold

PROMPT: which acquisitions need review - suggest anything with AI features, or spend above a threshold, whichever catches more.

Vendor questionnaire

Before purchase, the vendor must document: training-data provenance claims; where data is processed and whether it trains on your inputs; model updates and change notice; security certifications; subprocessors; how to export or delete your data.

Contract requirements

Data-processing agreement, input/output ownership, liability for AI-caused harm, notice of material model changes, exit and deletion terms.

Risk classification

The use case is classified against your risk framework (e.g. EU AI Act classes) BEFORE approval - high-risk uses need extra review and a named owner.

Register entry

No AI system goes live until it is entered in the AI system inventory with its owner and review date.

3. Human Oversight Policy

(c) Regovern.io - Free to use and adapt. More resources at https://regovern.io/resources/

Principle

Decisions with material effect on people's rights, livelihoods, safety, or access to services require meaningful human review before taking effect.

What 'meaningful' means

The reviewer has authority to overrule, understands what the system does and its known failure modes, has time to review properly, and is accountable for the outcome.

Oversight measures

PROMPT: per system - sampling rates, override logging, escalation paths, and what triggers a pause of automated operation.

Automation bias

Reviewers periodically see cases where the AI was wrong, and metrics on override rates are reported to the governance owner.

Records

Keep evidence of review decisions sufficient to demonstrate compliance (log excerpts, QA samples). PROMPT: retention period.

Making these yours

- Assign each policy a named owner and a review cycle (annual is typical).
- Version and date every policy; keep the previous version.
- Map each policy to the evidence an auditor would ask for (training records, exception logs).
- Keep the language plain: these documents should be readable by the people they bind.

(c) Regovern.io - Free to use and adapt. More resources at https://regovern.io/resources/