SAGAZ CONSULTING · GOVERNANCE TOOLS

AI Governance Policy Template

A board-ready starting policy for organizations that are already using AI and have not yet said, in writing, who decides what, what may be entered into which tool, and who is accountable for the output.

Prepared by Lcda. Raquel Núñez-Alicea · Attorney, Executive Advisor · hello@sagaz.consulting · sagaz.consulting

How to use this template

Replace every bracketed field with your own. The bracketed fields are deliberate: a policy without named owners, named tools and named timeframes is a statement of intent, not governance.

> **Before you adopt it:** this template is general information, not legal advice, and it does not create an attorney-client relationship. Sector rules (health, education, finance, government, regulated professions) and your own contracts may require more. Have it reviewed against your jurisdiction and your obligations before adoption.

The five fields that matter most

FIELD	WHAT TO PUT
[ROLE]	The single person accountable for AI governance. Not a committee.
[GOVERNING BODY]	Board, executive committee or owner who adopts the policy.
[CONFIDENTIAL TIER]	The classification above which approved tools are mandatory.
[ANNUALLY]	Your actual review cadence. Annual is a floor, not a target.
[24 HOURS]	Your incident reporting window.

Artificial Intelligence Use and Governance Policy

[ORGANIZATION] · Adopted [DATE] · Owner [ROLE] · Version 1.0

1. Purpose

This policy states how [ORGANIZATION] authorizes, governs and documents the use of artificial intelligence tools in its work. It exists so that decisions about AI are made by named people, under stated criteria, with a record that can be reviewed later.

2. Scope

This policy applies to all employees, officers, contractors, vendors and volunteers of [ORGANIZATION] who use, procure, build or supervise AI tools in connection with its work — including general-purpose assistants, embedded AI features inside approved software, and any model accessed through an API.

3. Definitions

AI tool — any software that generates text, code, images, audio, video, predictions, scores or recommendations using a machine-learning model.
Approved tool — an AI tool listed on the [ORGANIZATION] approved-tool register.
Confidential information — as defined in [ORGANIZATION]'s existing confidentiality, privacy and records policies; this policy does not replace those definitions.
Human review — review by a qualified person who is accountable for the output and has the authority and the time to change or reject it.

4. Principles

a. Accountability does not transfer to a tool. The person who signs, sends, files or publishes the work owns it.
b. AI may assist analysis and drafting. It does not make decisions reserved to professional judgment, management authority or the governing body.
c. Nothing enters an AI tool that could not lawfully be shared with the vendor behind it.
d. Where AI materially shapes a deliverable or a decision affecting a person, that is disclosed as required by section 8.
e. Use is documented well enough that a reviewer, auditor or regulator could reconstruct what happened.

5. Approved and prohibited uses

Permitted, with human review: drafting internal documents, summarizing material the organization already holds, research starting points, code assistance in non-production environments, brainstorming and reformatting.

Prohibited without written authorization from the [ROLE]: entering confidential, personal, health, financial, student, client or privileged information into a tool that is not approved; using AI output as the sole basis for a decision about a person (hiring, discipline, admission, credit, benefits, clinical or legal outcomes); generating content that is published or filed externally without named human review; building or deploying a model trained on organizational data.

6. Data handling

Before any information is entered into an AI tool, the user must classify it under [ORGANIZATION]'s existing data classification scheme. Information classified as [CONFIDENTIAL TIER] or above may only be entered into tools on the approved register that carry a written agreement covering confidentiality, data retention and the vendor's training rights.

The default assumption is that anything entered into a consumer AI tool leaves the organization's control.

7. Human review and accountability

Every deliverable with AI involvement has a named accountable person before it leaves the organization. Review is substantive: facts, citations, figures, names, legal and clinical statements are verified against sources. “It looked right” is not review.

8. Disclosure

[ORGANIZATION] discloses AI involvement where (a) the law, a regulator, an accreditor or a contract requires it; (b) the output materially affects a person's rights, treatment, grade, employment or money; or (c) a reasonable recipient would consider the involvement material. The [ROLE] maintains the standard disclosure language.

9. Tool approval

New AI tools are requested in writing and evaluated by the [ROLE] against: purpose, data to be processed, vendor terms (confidentiality, retention, training on inputs, subprocessors), security posture, cost, and the consequence of a wrong output. Approved tools, their approved uses and their approval date are recorded on the register. Approval is reviewed at least [ANNUALLY].

10. Training

Everyone in scope completes AI use and governance training before authorization, and a refresher at least [ANNUALLY]. Training covers this policy, the approved register, data handling, review duties, disclosure and incident reporting. Completion is recorded.

11. Incidents

A suspected incident — confidential information entered into an unapproved tool, a harmful or fabricated output acted upon, a vendor breach, an output that produced a wrong decision — is reported to [ROLE] within [24 HOURS]. Incidents are logged with date, tool, information involved, impact, corrective action and follow-up.

12. Records and evidence

[ORGANIZATION] maintains: the approved-tool register; written tool approvals; training completion records; the incident log; and, for decisions affecting individuals, a note of the AI involvement and the human review performed. Records are retained per [RETENTION POLICY].

13. Review of this policy

This policy is reviewed at least [ANNUALLY] by [ROLE] and approved by [GOVERNING BODY]. Material changes in law, tooling or organizational use trigger an earlier review.

14. Adoption

Adopted by [GOVERNING BODY] on [DATE]. Owner: [ROLE]. Next review: [DATE].

Decision owner matrix

Fill this in before the policy is adopted. If a row has no name, that decision is not governed.

DECISION	WHO DECIDES	EVIDENCE PRODUCED
Authorize a new AI tool	[ROLE]	Written approval on the register
Classify data before input	The user	Classification per existing scheme
Approve external AI-assisted deliverables	Named accountable person	Review recorded
Decide disclosure in a borderline case	[ROLE]	Decision noted in the file
Receive and log incidents	[ROLE]	Incident log entry
Approve and review this policy	[GOVERNING BODY]	Minutes and dated policy version

Where organizations get this wrong

A policy with no owner. Every section here assigns a person. If the answer is “the committee”, nothing is decided between meetings.

Prohibitions no one can follow. “Do not use AI” is not policy; it is a rule that gets broken quietly. Authorize specific tools for specific uses.

No evidence. If the organization cannot show the register, the training record and the incident log, it cannot demonstrate governance after something goes wrong.

© SAGAZ Consulting. You may adapt this template for your own organization's internal use. General information only — not legal advice.
Questions, or want it reviewed against your obligations? hello@sagaz.consulting · sagaz.consulting/ai-governance-session