AI Governance Framework for Wellness Organizations

A framework for building lightweight, enforceable AI governance. Includes policy templates, review cadences, and escalation paths.

Wellness organizations are adopting AI faster than they are governing it. A scheduling bot goes live before anyone writes down what it is allowed to touch. A staff member starts drafting client intake summaries with an AI tool before anyone decides what data is safe to paste into it. None of that is reckless. It is just unowned.

This framework builds governance that is light enough for a five-person studio and sturdy enough for a multi-location practice. It covers what to govern, the policies to write, how often to review them, and what to do when something goes wrong.

01 READINESS CHECK

Governance fails when it is built in the abstract. Confirm these four things before you write a single policy.

- ☐ You can list every AI tool currently in use, not just the ones you approved.
  Shadow AI, the tool staff started using without asking, is the most common gap. Ask before you assume.
- ☐ You have a named owner for AI governance, not a committee.
  One person accountable for keeping policies current beats a group that meets once a year.
- ☐ You know which data your organization cannot let an AI tool see.
  Client health information, session notes, and payment details need a hard line, not a judgment call.
- ☐ You have a way to log incidents when something goes wrong.
  A shared doc is enough to start. The point is that it exists before you need it.

02 WHAT TO GOVERN

AI touches a wellness organization in three distinct ways, and each one needs a different kind of check.

1. Agents that act on their own.
   A booking assistant, an automated intake bot, a follow-up messenger. Anything that takes action without a person clicking send needs a boundary on what it is allowed to do.
2. Work built with AI help.
   Staff using AI to draft client communications, protocols, or marketing copy. The output still has to meet your standards, even if AI wrote the first draft.
3. Any device that talks to an AI service.
   Every laptop, tablet, or phone a staff member uses to reach an AI tool is a place your data can leave the building. Govern the device, not just the tool.

Most organizations only think to govern the first. The other two are where the real exposure sits.

03 POLICY TEMPLATES

1. Acceptable Use Policy.
   Names which AI tools are approved, which are banned, and who to ask when a new one shows up.
2. Data Handling Policy.
   Defines exactly what client and business data can never be entered into an AI tool, with named categories, not vague warnings.
3. Human Override Policy.
   States which decisions an AI can never make alone: clinical judgment, refunds over a set amount, anything client-facing that involves health.
4. Vendor and Tool Approval Policy.
   A short checklist a new AI tool has to pass before staff can use it, covering data handling, security, and cost.
5. Incident Response Policy.
   Who gets notified, how fast, and what happens first when an AI tool does something it should not.

---

04 REVIEW CADENCE

A policy that is never revisited is a policy that is already out of date. Set the cadence now.

CADENCE	WHO	WHAT GETS REVIEWED
Weekly	Team lead	New AI tools staff have started using. Any incidents flagged that week.
Monthly	Governance owner	Policy exceptions granted, the incident log, and staff feedback on friction points.
Quarterly	Owner + leadership	Full policy review against actual day-to-day practice. Vendor and tool list audit.
Annually	Owner, plus outside advisor if needed	Complete framework rewrite. Regulatory changes. Insurance and liability review.

05 ESCALATION PATH

1. Tier 1 — Automated flag or staff report.
   Anyone can report a concern. It does not require proof, just a reasonable flag.
2. Tier 2 — Team lead review within 24 hours.
   Confirm what happened, contain it if needed, decide if it goes further.
3. Tier 3 — Governance owner review within 3 business days.
   Determine root cause, update the policy if the gap is systemic, log the resolution.

4. Tier 4 — Owner and outside counsel, only when required.
   Client data exposure, regulatory questions, or anything with legal exposure goes here. Most incidents never reach this tier.

06 WHERE THIS BREAKS

- Writing the policy once and never revisiting it.
  A policy that has not changed in a year usually has not been read in a year either.
- No consequence for bypassing an approved tool.
  If shadow AI use has no cost, staff will keep finding faster unapproved options.
- Treating every incident like it needs the owner's attention.
  Escalation tiers exist so small issues stay small. Route everything to the top and the real issues get lost in the noise.
- Governance that only covers agents, not staff-facing tools.
  The chat window a staff member opens on their own laptop is as much a governance gap as any automated bot.
- No log, so no way to spot a pattern.
  One incident is a fluke. Three of the same incident is a policy that is not working.

Governance does not slow a wellness organization down. Unowned AI use does, eventually, all at once.

The organizations that get this right do not have the most restrictive policies. They have the clearest ones, owned by one person, reviewed on a schedule, with an escalation path everyone already knows before they need it.