KICKSTANDAI x MY-CC.IO

AI Governance and Compliance

A practical framework for building responsible, defensible AI oversight in healthcare and wellness organizations.

EXECUTIVE SUMMARY

AI adoption in healthcare and wellness organizations has outpaced the policies meant to govern it. Scheduling agents, intake automation, content generation, and staff-facing AI tools are now standard operating infrastructure, often introduced faster than any formal review process can keep up with. AI governance is the set of policies, processes, and controls that ensure these tools are used responsibly, legally, and in line with organizational values. Compliance is the narrower discipline of meeting the specific regulations that apply to your industry, your data, and your region. The two are related but not identical: an organization can be technically compliant and still ungoverned, or governed and still exposed to regulatory risk. This white paper lays out the core components of an AI governance program, the compliance considerations specific to healthcare and wellness organizations, and a practical, phased path to building a program that is sustainable rather than performative.

IN THIS WHITE PAPER

01 What Is AI Governance and Compliance?
02 The Core Components of AI Governance
03 Compliance Considerations by Sector
04 Building a Practical Governance Program
05 Common Pitfalls
06 A 90-Day Implementation Timeline
07 Key Takeaways

01 WHAT IS AI GOVERNANCE AND COMPLIANCE?

AI governance is the set of policies, processes, and controls that ensure AI tools are used responsibly, legally, and in line with organizational values. It answers a specific set of questions: which tools are approved, who can use them, what data they can touch, and what happens when something goes wrong. Governance is not a single document. It is a system of recurring decisions, applied consistently, that holds up whether the person enforcing it is in the room or not.

Compliance means meeting the regulations that apply to your industry and region. For a healthcare organization, that includes HIPAA and state-level privacy statutes. For a wellness or plant-based consumer brand, it includes FTC guidance on marketing claims, state consumer protection law, and data retention requirements tied to e-commerce and loyalty programs. Compliance is externally defined: a regulator, an auditor, or a plaintiff's attorney decides whether you met the bar, not your own internal comfort level.

The distinction matters because organizations often build one without the other. A practice can pass a HIPAA risk assessment and still have no idea which AI tools its front desk staff downloaded last month. A wellness brand can have an airtight privacy policy on its website and still let an AI content tool generate a health claim no one reviewed before it went out. Governance without compliance leaves you well-organized and still exposed. Compliance without governance leaves you defensible on paper and unaware of what is actually happening day to day.

A functioning program treats governance as the operating system and compliance as one of the things that operating system has to run correctly. Build governance first, and compliance becomes a natural output of the review cycles, access controls, and oversight you already have in place, rather than a separate project bolted on after the fact.

---

02 THE CORE COMPONENTS OF AI GOVERNANCE

1. **Policy**

   Policy is the documented rule for when and how AI may be used. It should name the approved tools, the tools that are explicitly banned, and the process for evaluating a new one. A policy that lives only as an unwritten understanding among senior staff is not a policy. It is a habit that leaves with whoever holds it in their head. Write it down, date it, and assign someone to keep it current.

2. **Access Control**

   Access control determines who can use which tools and what data those tools can reach. Not every staff member needs the same level of AI access, and not every AI tool needs the same level of data access. A content-drafting tool used by a marketing coordinator has no legitimate reason to touch a client's clinical record. Scope access to the job, not to the org chart.

3. **Review Cycles**

   Review cycles are the regular audits of agent behavior, outputs, and permissions. A tool approved six months ago under one set of assumptions may be operating under very different ones today, especially if it has been updated by the vendor or connected to a new data source. Reviews catch that drift before it becomes an incident. They also catch the more common problem: tools nobody remembers approving in the first place.

4. **Human Oversight**

   Human oversight is the set of escalation paths and approval requirements for high-stakes actions. Not every AI-assisted decision needs a human in the loop, but some categorically do: clinical judgment, financial decisions above a set threshold, anything client-facing that touches health status. Oversight is not a sign that you do not trust the tool. It is an acknowledgment that some decisions carry consequences a policy document cannot absorb on its own.

---

03 COMPLIANCE CONSIDERATIONS BY SECTOR

The regulatory floor is different depending on what you sell and who you serve. Two organizations running the exact same AI governance program can have very different compliance obligations layered on top of it.

Healthcare Organizations

Healthcare organizations operate under HIPAA at the federal level, plus a patchwork of state privacy laws that vary in scope and enforcement. Professional licensing boards add another layer: a clinician using AI to draft a treatment note is still fully responsible for that note's accuracy and appropriateness, regardless of what tool helped write it. AI governance in this context has to account for protected health information specifically, not data in general. A tool that is safe for scheduling data can be entirely unsafe for a clinical note, even if it is the same tool used by the same staff member five minutes later.

Plant-Powered and Wellness Brands

Plant-based and wellness consumer brands face a different set of obligations. Consumer protection law and FTC guidance govern what you can claim about a product's health effects, and AI-generated marketing copy is not exempt just because a person did not type it themselves. Data retention requirements apply to customer purchase history, loyalty program data, and anything collected through a coaching or intake form. The reputational risk runs alongside the regulatory one: a customer base that chose you for your values will notice an AI-generated claim that oversells what your product does, long before a regulator does.

SECTOR	PRIMARY REGULATORY CONCERN	HIGHEST-RISK AI USE CASE
Healthcare	HIPAA, state privacy law, professional licensing	Clinical documentation or patient communication drafted or summarized by AI.
Wellness / Plant-Based Brands	FTC marketing guidance, consumer protection law, data retention rules	AI-generated health or efficacy claims in marketing content.
Both	Data collected through automated intake or coaching workflows	Any AI tool with standing access to personal data across multiple sessions.

Organizations that serve both markets, a wellness brand with a clinical coaching arm, for example, inherit the higher bar. Build the program to the stricter standard and the lighter one is already covered.

---

04 BUILDING A PRACTICAL GOVERNANCE PROGRAM

1. **Start with an AI use inventory.**
   List every tool currently in use, who uses it, what data it touches, and what decisions it influences. This step alone typically surfaces more shadow AI than anyone expects. Do not skip it in favor of writing policy first. A policy built without an accurate inventory is a policy built for a workplace that does not exist.

2. **Assign a risk rating to each tool.**
   Rate every tool on the inventory by what it touches and what it decides, not by how impressive the underlying technology is. A tool with access to health data or financial decisions is high risk even if it performs a simple task. A tool that only drafts internal meeting notes is low risk even if it is technically sophisticated. Risk follows exposure, not capability.

3. **Set a review schedule proportional to the risk.**
   High-risk tools get reviewed monthly. Low-risk tools can go quarterly. This is where most governance programs either become sustainable or quietly stop happening. A review cycle that is too aggressive for a small team will get skipped the first time things get busy. A review cycle that is too loose for a high-risk tool will miss the incident that actually matters.

4. **Formalize the escalation path before you need it.**
   Decide, in writing, what happens when a tool does something it should not: who gets notified, how fast, and what the first containment step is. An escalation path decided in the moment is usually the wrong one, made under pressure, by whoever happened to notice the problem first.

5. **Revisit the whole program on a fixed cadence, not an ad hoc one.**
   Governance is not a project with an end date. Set a recurring date, quarterly at minimum, to review the inventory, the risk ratings, and the policy itself against what is actually happening in the organization. Programs that are reviewed only when something goes wrong are not governance. They are damage control with a policy binder attached.

05 COMMON PITFALLS

— Writing governance policy before completing the AI use inventory.
A policy built on assumptions about tool use, rather than a documented inventory, tends to miss the tools staff actually rely on.

— Rating every tool the same, regardless of what it touches.
Treating a scheduling reminder tool the same as a clinical documentation tool wastes review time on the former and under-reviews the latter.

— Assuming compliance covers governance.
Passing an audit tells you what a regulator checked. It does not tell you what your staff started using last week.

— No consequence for bypassing an approved tool list.
A policy with no enforcement is a suggestion, and staff under time pressure will treat it like one.

— Reviewing the program only after an incident.
By the time a review is triggered by a failure, the cost of that failure has already been paid.

06 A 90-DAY IMPLEMENTATION TIMELINE

A governance program built all at once, everywhere, is a governance program nobody follows. Stage it.

PHASE	FOCUS	WHAT SHIPS
Days 1–30	Discovery	AI use inventory completed. Initial risk ratings assigned. Governance owner named.
Days 31–60	Formalize	Core policies drafted and circulated: acceptable use, data handling, human override. Review cadence set per risk tier.

Days 61–90	Operationalize	First full review cycle run. Escalation path tested with a tabletop exercise. Lessons learned documented and folded back into policy.

07 KEY TAKEAWAYS

- **Governance is about repeatable decisions, not one-time documents.**
  A policy that is not applied consistently, month over month, is not really governance. It is paperwork.
- **Compliance starts with knowing which rules apply to your data and workflows.**
  You cannot comply with a regulation you have not identified as relevant to what you actually do.
- **Human oversight is the strongest control you can add.**
  Every other control, policy, access rules, review cycles, exists to make sure a human is positioned to catch what automation misses.

A governance program is not a compliance shield. It is how an organization stays honest about what its AI tools are actually doing.

Start with the inventory. Rate what you find honestly. Review it on a schedule you will actually keep. Everything else in this white paper is a refinement of that basic loop, repeated consistently until it is simply how the organization operates.
