

Responsible AI Use: A Practical Framework

Simple, durable guardrails for using AI in regulated and high-stakes work

Adopting AI responsibly does not require a 40-page policy. It requires a few clear principles everyone understands and follows. This framework is drawn from using AI daily inside a HIPAA-regulated healthcare environment, where the cost of getting it wrong is real.

The five principles

1. AI output is a draft, never final truth

Every AI output is a starting point that a human, and where needed a subject-matter expert, verifies before it is used or shared. This is the single most important rule. It protects accuracy, compliance, and credibility.

2. Protect the data

Define in plain language what information may and may not be entered into a given AI tool. Sensitive, regulated, or confidential data goes only into approved, compliant systems, never into general consumer tools. When in doubt, leave it out.

3. Keep a human accountable

AI assists; a named person remains responsible for the result. Automation never removes ownership. This keeps decision-making and accountability clearly human.

4. Be transparent about use

Establish a team norm for when and how AI assistance is disclosed. Transparency builds trust internally and with the people you serve and prevents quiet over-reliance.

5. Match the tool to the stakes

Low-stakes, reversible tasks (drafting, brainstorming) are ideal for fast AI assistance. High-stakes, hard-to-reverse, or compliance-sensitive work demands heavier verification and, often, a human-led process with AI only in a supporting role.

Why this matters for enablement

Teams adopt AI faster and more sustainably when the guardrails are clear, because clarity removes fear. People who know exactly where the lines are will use AI more confidently, not less. Good governance is what makes adoption safe enough to scale.