

# The AI Readiness Maturity Model

For health and wellness practices.

---

V1.0 · SHAUN DUVAL · FRACTIONAL CHIEF AI OFFICER

---

## What this is

A self-scoring assessment for SMB health and wellness practices (med spas, group therapy practices, holistic practitioners, wellness brands, multi-location operators) to figure out where their AI posture actually sits, and what the next move is.

6 pillars. 5 levels each. 24 questions, 12 minutes.

The output is a per-pillar score and a written next step. The framework is the same one I use in the paid Operating Assessment, stripped down to what you can run yourself.

## How to use it

1. Answer each question honestly. Pick the option that's closest to true today, not the one closest to your ambition. Aspiration is not readiness.
2. Each answer maps to a level, 1 to 5.
3. Average the answers inside a pillar to get that pillar's level.
4. Read the "what to do next" notes for your weakest pillar first.

Score the floor, not the ceiling. The number matters less than the gap it surfaces.

---

## The 6 pillars

1. **Governance.** The written rules: who decides, what's allowed, where the legal lines are (HIPAA, state board, FTC, ADA).
2. **Tools and Data.** The AI tools in use, who picked them, and what data they touch.
3. **Voice.** Whether the practice's voice survives contact with AI output, or gets eaten by generic model writing.
4. **Accessibility.** Whether the things you ship work for people using screen readers, keyboard navigation, captions, or plain language.
5. **People.** Staff AI literacy, training cadence, and who actually owns the work day to day.
6. **Measurement.** Whether AI is moving a KPI you can name, or just running in the background while you hope it's working.

## The 5 levels

Each pillar uses the same 5-level scale.

1. **Unmapped.** Nobody owns this. The question hasn't been asked.
2. **Ad-hoc.** Decisions get made case by case, in DMs and Slack threads, with no written record.
3. **Defined.** A written policy or practice exists. People can point to it. Adoption is uneven.
4. **Operating.** The practice is running, being followed, and getting reviewed on a cadence.
5. **Compounding.** Measured, improved over time, and producing a real advantage (faster onboarding, fewer legal questions, lower complaint rate, defensible voice).

A practice can be Operating on Tools and Unmapped on Accessibility at the same time. The point of the framework is to surface that asymmetry.

---

---

## Pillar 1: Governance

The written rules. Who is allowed to use which AI tool, for what, on what data, with what disclosure to patients or clients. The HIPAA Business Associate boundary. The FTC substantiation line for marketing claims. The state board posture for anything that touches clinical delegation.

### WHAT EACH LEVEL LOOKS LIKE

- **1. Unmapped.** No AI policy. Staff use whatever they downloaded. Nobody at the practice can name the BAA boundary.
- **2. Ad-hoc.** The owner has a vague rule ("don't put patient names in ChatGPT") but it's not written down. New hires never hear it.
- **3. Defined.** A written Acceptable Use Policy exists. It names approved tools, prohibits PHI in non-BAA tools, and covers FTC substantiation for AI-generated marketing.
- **4. Operating.** The AUP is signed at onboarding, reviewed annually, and tied to vendor selection. Someone owns it.
- **5. Compounding.** Governance is a live document. Incidents and near-misses get logged and feed back into the policy. Vendors are scored against a written rubric before adoption.

### QUESTIONS

#### Q1. Does your practice have a written AI Acceptable Use Policy?

1. No, and we haven't discussed it.
2. We have an informal rule but nothing written.
3. Yes, it exists in a shared drive.
4. Yes, every staff member signs it at onboarding.
5. Yes, and we review and update it at least annually.

**Q2. If a staff member wants to try a new AI tool tomorrow, what happens?**

1. They install it and start using it.
2. They mention it to a manager who decides on the spot.
3. There's a process but most people skip it.
4. They submit it for review against a written checklist.
5. They submit it, it gets scored against a rubric, and the decision is logged.

**Q3. Can you name the HIPAA Business Associate boundary for your practice in one sentence?**

1. No.
2. I know what HIPAA is, not the BA piece.
3. Roughly, but I'd have to look it up.
4. Yes, and it's documented.
5. Yes, it's documented, and every vendor is mapped against it.

**Q4. If a regulator (OCR, FTC, state board) asked you to produce your AI usage records this week, what would you send?**

1. Nothing exists.
  2. Screenshots and a Slack export.
  3. A policy document.
  4. A policy, vendor inventory, and BAAs on file.
  5. All of the above plus an incident log and review history.
-

---

## Pillar 2: Tools and Data

The actual AI in use. Which models, which vendors, what data flows through them. Whether anything is hitting protected health information without a Business Associate Agreement in place. Whether the practice could answer "what AI tools are we paying for?" without opening QuickBooks.

### WHAT EACH LEVEL LOOKS LIKE

- **1. Unmapped.** Nobody can list the AI tools the team uses. PHI may be passing through consumer ChatGPT and nobody would notice.
- **2. Ad-hoc.** A few tools are known. A few are not. BAAs exist for some, not others.
- **3. Defined.** A current inventory of AI tools exists. Each tool is tagged with what data it touches and whether a BAA is in place.
- **4. Operating.** The inventory is current within 30 days. Off-list usage is flagged. New tools go through procurement.
- **5. Compounding.** The inventory drives renewals, cost reviews, and consolidation. Vendor risk is scored and re-scored.

### QUESTIONS

**Q5. Can you produce a current list of every AI tool your team uses for work?**

1. No.
2. A partial list, missing the long tail.
3. Yes, but it's a few months stale.
4. Yes, updated within the last 30 days.
5. Yes, with vendor risk scoring attached.

**Q6. For each AI tool that could touch patient or client data, is there a Business Associate Agreement on file?**

1. We don't know which ones touch patient data.
2. Some have BAAs; we're not sure which.
3. Most do. The list isn't authoritative.
4. Yes, mapped one-to-one with the tool inventory.
5. Yes, with renewal dates tracked.

**Q7. If a staff member pasted a patient name into ChatGPT today, how would you find out?**

1. We wouldn't.
2. Eventually, if they mentioned it.
3. Through a manual review.
4. Through a documented monitoring process.
5. Through monitoring and a no-blame reporting channel.

**Q8. How are AI vendor decisions made today?**

1. Whoever signed up first.
  2. The owner picks based on a demo.
  3. A short internal discussion.
  4. A written evaluation against criteria.
  5. A scored rubric with named owners and renewal triggers.
-

---

## Pillar 3: Voice

Whether what AI produces still sounds like the practice. The intake script, the welcome email, the social caption, the consult summary, the bio paragraph. Founders spend years building a voice. AI eats that voice in week 3 unless something protects it.

### WHAT EACH LEVEL LOOKS LIKE

- **1. Unmapped.** AI output goes out unedited. The voice is already drifting and nobody on staff has named the drift.
- **2. Ad-hoc.** One person edits when they catch it. The catch rate is low.
- **3. Defined.** A written voice guide or style brief exists. Editors reference it inconsistently.
- **4. Operating.** Every AI-generated piece passes a voice check before it ships. The voice guide is loaded into the tools.
- **5. Compounding.** The voice guide is versioned, the tools are fine-tuned or prompted against it, and new staff are calibrated with examples.

### QUESTIONS

**Q9. Is there a written voice guide or style brief for your practice?**

1. No.
2. Some scattered notes.
3. Yes, in a shared drive.
4. Yes, and it's loaded into the AI tools that write.
5. Yes, versioned, with examples of in-voice and out-of-voice.

**Q10. Pull the last 3 things AI helped your team write. How many sound like the practice?**

1. Zero.
2. One.
3. Two.
4. All three, after voice editing.
5. All three, with no voice editing needed.

**Q11. When new staff write the first AI-assisted piece, what happens?**

1. It ships.
2. Somebody catches it later.
3. A manager reviews it.
4. It goes through a voice check against the guide.
5. It goes through a voice check and the result feeds the guide.

**Q12. Pick a piece of AI-assisted content you shipped this week. Does it sound like you?**

1. I'd be embarrassed if a longtime client saw it.
  2. It sounds like AI writing about my topic.
  3. It sounds like me on a bad day.
  4. It sounds like me.
  5. It sounds like me at my best.
-

---

## Pillar 4: Accessibility

Whether the things AI helps you ship actually work for people using screen readers, keyboard-only navigation, captions, or plain language. ADA, Section 508, and the European Accessibility Act all apply to health and wellness, and the exposure is real. AI tools generate inaccessible defaults unless someone catches them.

### WHAT EACH LEVEL LOOKS LIKE

- **1. Unmapped.** Accessibility hasn't been discussed. The site, forms, and patient portal have never been tested.
- **2. Ad-hoc.** Someone added alt text to a few images once.
- **3. Defined.** A written accessibility standard exists (WCAG 2.2 AA, typically). Reviews happen on request.
- **4. Operating.** Every shipped surface (site, forms, patient portal, video content) goes through an accessibility check before launch.
- **5. Compounding.** Accessibility is in the brief from the start, not the QA at the end. Real users on assistive tech are part of testing.

### QUESTIONS

**Q13. When was the last time someone ran an accessibility check on your website or patient portal?**

1. Never.
2. Once, a while back.
3. Within the last year.
4. Within the last 90 days.
5. As part of every release.

**Q14. Do AI-generated images on your site have meaningful alt text?**

1. No alt text at all.
2. Alt text exists but it's auto-generated noise.
3. Alt text exists and is mostly accurate.
4. Alt text is written or reviewed by a human against a guide.
5. Alt text is reviewed and the guide updates with new patterns.

**Q15. If a patient using a screen reader tried to book an appointment today, would they succeed?**

1. We don't know and we've never tested.
2. We assume so.
3. We tested once and it mostly worked.
4. We test every release.
5. We test every release with an actual screen reader user in the loop.

**Q16. Are videos and audio content captioned and transcribed?**

1. No.
  2. Some, when we remember.
  3. Yes, with auto-generated captions only.
  4. Yes, with reviewed captions and transcripts.
  5. Yes, plus alt audio descriptions where the visual matters.
-

---

## Pillar 5: People

Whether the staff using AI actually know how to use it, how to escalate when something looks off, and who owns the AI work internally. Tools without trained operators are the most common failure I see.

### WHAT EACH LEVEL LOOKS LIKE

- **1. Unmapped.** No training. Staff figured it out on YouTube.
- **2. Ad-hoc.** Some staff are good at AI, some are not. There's no pattern.
- **3. Defined.** A training curriculum exists. New hires get a walkthrough.
- **4. Operating.** Training happens on a cadence. There's a named internal AI lead.
- **5. Compounding.** Training is updated as tools change. The AI lead runs office hours. Staff bring back use cases that improve the policy.

### QUESTIONS

**Q17. Does any staff member have written training on the AI tools they use?**

1. No.
2. Informal coaching from a peer.
3. A one-time training session.
4. Documented training they refer back to.
5. Documented training plus a refresh cadence.

**Q18. Is there one named person responsible for AI inside the practice?**

1. No.
2. The owner, by default.
3. A staff member volunteered.
4. A named lead with allocated time.
5. A named lead with time, budget, and a written remit.

**Q19. If a staff member sees an AI tool make a mistake (wrong patient info, off-brand copy, hallucinated claim), what do they do?**

1. Ignore it or fix it silently.
2. Tell a manager when they remember.
3. Report it through a known channel.
4. Report it through a logged channel reviewed weekly.
5. Report it, it's logged, reviewed, and feeds the policy.

**Q20. How often does the team talk about AI in a structured way (not in passing)?**

1. Never.
  2. When a problem comes up.
  3. Quarterly.
  4. Monthly.
  5. Weekly or biweekly office hours.
-

---

## Pillar 6: Measurement

Whether AI is moving a number you can defend. Most practices buy AI tools, run them for a year, and never check whether anything got better. "We save time" is not a KPI. The retainer-tier work I do hinges on naming one KPI per practice and tracking it monthly. Practices that can't name one are guessing.

### WHAT EACH LEVEL LOOKS LIKE

- **1. Unmapped.** No KPIs tied to AI. Nobody could tell you whether AI is helping or just adding cost.
- **2. Ad-hoc.** Anecdotes ("front desk likes it", "feels faster") but no numbers.
- **3. Defined.** One or two KPIs named, baseline captured. Reporting is irregular.
- **4. Operating.** Named KPIs reviewed on a cadence (monthly or quarterly). Trends are tracked against the baseline.
- **5. Compounding.** New AI investments require a written hypothesis, a baseline, a target, and a review cadence before rollout. Decisions to keep or kill tools are logged.

### QUESTIONS

**Q21. Can you name a KPI that AI is moving in your practice today?**

1. No.
2. Only in general terms (we save time).
3. Yes, one KPI.
4. Yes, with a current value and a baseline.
5. Yes, with current value, baseline, and a target.

**Q22. When you adopted your last AI tool, what was the measurement plan?**

1. There wasn't one.
2. We figured we'd notice if it helped.
3. We agreed informally on what to watch.
4. We named a KPI and a baseline before rollout.
5. We named the KPI, baseline, target, and review cadence.

**Q23. How often do you review whether AI is hitting outcomes (not whether the tools are running, but whether they're moving numbers)?**

1. Never.
2. When something breaks.
3. Quarterly.
4. Monthly with a named owner.
5. Monthly, with a written report to leadership.

**Q24. If I asked your team "is the AI working?" today, what would they show me?**

1. A guess.
  2. A demo of the tool.
  3. A few anecdotes.
  4. A KPI trend.
  5. A KPI trend with attribution and a written decision history.
-

# Scoring

For each pillar:

- Add up the answers (each answer = its option number, 1 to 5).
- Divide by 4 to get the pillar level.

Pillar levels are usually a decimal (e.g. 2.3). Round down. A 2.9 rounds to 2. Aspiration is not readiness.

Then add the 6 pillar scores together for an overall score out of 30.

## WHAT THE OVERALL SCORE MEANS

- **6 to 11: Unmapped.** AI is happening in your practice whether you like it or not. The first move is naming it, not optimizing it. Start with Governance and Tools.
- **12 to 17: Ad-hoc.** You have momentum and no scaffolding. The fastest gain is writing down what already works and killing what doesn't. Start with Governance and the weakest of Voice or Accessibility.
- **18 to 22: Defined.** The policies exist. The next move is making them real. Add review cadences and Measurement to whichever pillar is lowest.
- **23 to 27: Operating.** You're running it. Now make it compound. Tighten Measurement on the pillar you can least defend, and start feeding incidents back into policy.
- **28 to 30: Compounding.** Rare. If this is honest, you're ahead of every solo health and wellness practice I see and most multi-location groups. Use the readiness as a moat.

## PER-PILLAR NEXT STEPS

Find your weakest pillar. Start there.

**Governance.** Write the Acceptable Use Policy this week. 1 page. Approved tools, prohibited data, BAA boundary, FTC line, who decides on new tools. Make new hires sign it.

**Tools and Data.** Build the inventory this month. Every AI tool the team touches, what data it sees, whether a BAA is in place. Cancel anything you can't justify.

**Voice.** Write the voice guide this quarter. 3 pages. Examples of in-voice and out-of-voice. Load it into every tool that writes for you.

**Accessibility.** Run a WCAG 2.2 AA audit on the site, patient portal, and intake forms this quarter. Fix the highest-severity issues. Bake the standard into every future release.

**People.** Name an internal AI lead. Give them 4 hours a week and a budget. Run office hours monthly. Track what the team brings back.

**Measurement.** Pick one KPI per active AI tool this month. Baseline it. Review it on the same cadence as the rest of your operating dashboard. Don't add a new AI tool without a written hypothesis and a measurement plan.

---

## What this assessment doesn't tell you

It doesn't tell you which AI tools to buy. It doesn't tell you what your roadmap should be. It doesn't tell you whether a specific vendor is right for your practice or whether your existing AUP holds up against a real OCR audit.

That's the Operating Assessment (\$9,500, 3 weeks): a full per-pillar diagnostic with vendor inventory, governance review, prioritized use cases, costed roadmap, and an executive readout.

Book a discovery call at [shaundual.com](https://shaundual.com).

---

*v1.0. The framework will be updated as the regulatory clocks move (OCR Security Rule, FTC Operation AI Comply, EU AI Act, state attorney general activity).  
Subscribers get the new version when it ships.*