



# The AI-Ready Counselor

A PRACTICAL WORKBOOK FOR MENTAL HEALTH PROFESSIONALS

---

*Preparing your practice for AI in payer systems, clinical work, and the business of therapy.*

Clarity. Compliance. Competitive advantage —  
without the hype and without the overwhelm.

## WELCOME

# Start here.

If you picked up this workbook, you're probably somewhere on a spectrum that runs from **curious** ("Should I try one of these AI note tools?") to **uneasy** ("Are insurance companies using AI to deny my claims?") to **quietly alarmed** ("If I don't do something soon, am I going to get left behind?").

All three responses are reasonable. AI is moving faster than most professions can comfortably absorb, and counseling is not exempt. Payers are already piloting AI in utilization review and claims processing. Clinical AI tools — scribes, note-drafters, screening assistants — are multiplying each month. And on the business side, AI is quietly becoming a default in marketing, scheduling, and admin work.

This workbook is not a hype piece, and it's not a doomsday warning. It's a structured, unhurried walk through the three areas that matter most to a private practice or group practice clinician right now:

## 01

### Payers & AI

How insurers are using AI in utilization review, claims, and prior auth — and how to write documentation that holds up to algorithmic review.

## 02

### Clinical AI

How to evaluate ambient scribes, note-drafters, and decision-support tools through the lens of HIPAA, ethics, and clinical quality.

## 03

### Business AI

Where AI adds real leverage in marketing, intake, scheduling, and admin — and where it creates risk you shouldn't accept.

### How to use this workbook

Work through it at your own pace. Each part ends with a worksheet you can fill in by hand or on screen. The final checklist (60+ items) is designed to be revisited quarterly as tools and regulations evolve. A single careful pass is worth more than a rushed one. Many clinicians find Part 1 (payers) the most immediately useful — feel free to start there.

**CONTENTS**

---

# What's inside

**01 The AI Moment in Counseling**

A plain-English map of what's actually happening.

---

**02 AI & Your Payers**

Utilization review, claims, prior auth, and the new importance of medical necessity language.

---

**03 AI in Clinical Practice**

Ambient scribes, note-drafters, screeners — benefits, risks, and the questions to ask every vendor.

---

**04 AI in Business Practice**

Marketing, SEO, intake, scheduling, admin — where the leverage is and where the line is.

---

**05 Compliance, Ethics & Informed Consent**

HIPAA, BAAs, licensing boards, ethics codes, and disclosure.

---

**06 Your AI Readiness Plan**

A structured worksheet: assess, prioritize, pilot, document, review.

---

**07 The Readiness Checklist**

60+ practical items, organized by domain.

---

*Estimated time: 2–4 hours spread over a week. Don't try to do it in one sitting.*

## PART 01

# The AI Moment in Counseling

*A plain-English map of what's actually happening.*

## The three-front reality

Most of what gets written about AI in mental health is either sales copy or alarm. Both obscure the actual picture, which is simpler and more useful: AI is arriving in your practice on three fronts at once, and each front moves at a different speed and carries different risks.

### The payer front (moving fastest — mostly invisible to you)

Insurance companies have been using algorithmic tools in claims and utilization review for years. What's new is that these tools are getting substantially more capable and are being deployed in more decisions — including first-pass medical necessity reviews, pattern-of-care analysis, and prior authorization. You rarely see it happen; you only see the result (an approval, a denial, a records request, a clawback letter).

### The clinical front (moving fast — highly visible)

Ambient AI scribes that generate progress notes from session audio, note-drafting tools that expand your bullet points into paragraphs, risk-screening add-ons — these are being marketed directly to clinicians and practice owners. The pitch is real (reduced admin burden, more eye contact), and so are the risks (HIPAA, accuracy, client consent, loss of clinical voice).

### The business front (moving fastest of all — and easiest to misuse)

General-purpose tools like ChatGPT, Claude, and Gemini are already embedded in how many solo and group practices handle marketing copy, blog posts, social media, intake templates, and admin email drafts. This is mostly upside — *until* PHI ends up in a public tool, or until a state licensing board takes a position on AI-generated marketing claims.

#### One principle to carry through the whole workbook

AI is a **force multiplier** — it amplifies both your strengths and your weaknesses. If your documentation is already tight, AI tools make it tighter. If it's loose, AI won't fix it — and payer AI will find the gaps faster than a human reviewer would. *Start with your fundamentals; then add tools.*

## Where are you right now?

Before going further, take a moment to locate yourself. Circle the phrase that best describes where your practice is today:



STAGE	DESCRIPTION
<b>Unaware</b>	Haven't looked at AI tools; no plan; no stated position with clients.
<b>Curious</b>	Experimenting personally; nothing in place at the practice level.
<b>Piloting</b>	Using 1–2 AI tools; informal consent; no formal vendor review.
<b>Structured</b>	Documented policy, BAAs signed, client consent process, limited use.
<b>Integrated</b>	AI is part of workflow; reviewed quarterly; clear client disclosures.

**Write your current stage and one sentence on why:**

---

---

---

## PART 02

# AI & Your Payers

*What's changing in utilization review, claims, and prior auth.*

## What payers are actually doing with AI

You don't need to be an insurance industry analyst to prepare for AI on the payer side. You only need a working mental model of the decisions where AI is now likely involved, and how to write documentation that survives algorithmic review.

### Decisions where payer AI shows up

Across commercial, Medicare Advantage, and Medicaid managed care plans, expect AI or AI-assisted algorithms to be present in some combination of:

<b>Claims adjudication</b>	First-pass review of submitted claims against the plan's rules, CPT/ICD combinations, place-of-service, and frequency norms.
<b>Utilization review</b>	Evaluating whether ongoing care meets medical necessity criteria, often by scanning submitted notes for specific markers.
<b>Prior authorization</b>	Pre-approving (or denying) a requested level of care based on submitted clinical information and guidelines.
<b>Outlier detection</b>	Flagging providers whose billing, session frequency, or diagnosis patterns differ meaningfully from peers.
<b>Fraud &amp; abuse screening</b>	Pattern analysis across large provider datasets.
<b>Records requests</b>	Automated identification of claims to pull for post-pay review (which can lead to clawbacks if documentation doesn't support what was billed).

### A shift worth internalizing

Human reviewers skim. **Algorithmic reviewers don't.** They check every field, every code, every date, every phrase against a rubric. That means small inconsistencies you used to get away with — a missing symptom statement, a vague goal, a session start time that doesn't match the CPT code — now produce actual denials.

## Writing documentation that holds up to AI review

The good news: writing for an AI reviewer is mostly the same as writing for a careful human reviewer. It's just less forgiving of sloppiness. Four habits matter more than any tool you could buy.

### 1. Medical necessity in every note

Every progress note should make the medical-necessity case implicitly: presenting symptoms with specificity, functional impairment, the intervention used, and the client's response. "Supportive listening; client feels better" is not a note that survives scrutiny. *"Client presented with persistent depressed mood (6/10), reduced appetite, and impaired concentration at work. Used cognitive restructuring targeting catastrophic thoughts about job performance; client identified two distortions and generated alternative appraisals."* is.

### 2. Goals that are measurable and tied to notes

Treatment plan goals should be observable and measurable, and each progress note should reference movement toward at least one goal. Algorithms look for the linkage between the plan and the session. When it's absent in note after note, that's a pattern — and patterns get flagged.

### 3. Consistent CPT, time, and documentation alignment

If you bill 90837 (53+ minutes of psychotherapy), the note should reflect actual start and end times that support that code. If you bill 90834, the same applies. Inconsistencies between billed time and documented time are one of the most common auto-flagged patterns.

### 4. Updated diagnoses and justified changes

When a diagnosis changes, the note should show why. When a diagnosis stays the same for many months, the notes should show that the symptoms still meet criteria. "By report, still meets criteria for MDD" is weaker than a brief criteria check written into the note every few months.

#### Field note

Many clinicians who get records-request letters describe the same experience: "My work was fine, but my notes didn't tell the whole story of my work." The fix isn't writing longer notes — it's writing notes that include the four elements above consistently.

## WORKSHEET · 01

# Your documentation audit

Pull **five of your recent progress notes at random**. For each, answer these four questions honestly. You are not grading yourself — you are building a baseline.

Check	Note 1	Note 2	Note 3	Note 4	Note 5
Specific symptoms with severity / impairment?					
Specific intervention named (not just "supportive")?					
Client response described?					
Links to at least one treatment plan goal?					
CPT code matches documented session length?					
Diagnosis still supported by content?					

**Pattern you noticed across the five notes:**

---



---



---

**One change you'll make to every note going forward:**

---



---



---

## PART 03

# AI in Clinical Practice

*Scribes, note-drafters, screeners — benefits, risks, questions.*

## A taxonomy of clinical AI tools

Vendors use overlapping language, which makes comparison hard. The tools actually sitting in front of you fall into five categories:

<b>Ambient AI scribes</b>	Record the session (audio or transcript), then produce a draft progress note. Largest time savings; largest privacy surface; consent conversations matter most here.
<b>Note-drafters</b>	You type bullets or a short summary; the tool expands into a full note. Less invasive than ambient scribes; still PHI if you include client details.
<b>Screeners / assessment helpers</b>	Administer or score measures (PHQ-9, GAD-7, etc.), flag risk indicators. Useful, but the clinician's judgment, not the score, remains the clinical decision.
<b>Clinical decision support</b>	Suggests interventions, differential diagnoses, or content. Treat output as a <i>second opinion to evaluate</i> , not a directive to follow.
<b>Between-session tools</b>	Chatbots or companion apps clients use between sessions. Highest ethical complexity; handle with explicit consent and scope-of-practice awareness.

### The benefits, honestly stated

Clinicians who use AI scribes responsibly report three consistent gains: **less after-hours documentation** (often the single biggest quality-of-life change), **more present attention** in the room (no laptop screen between you and the client), and **more thorough notes** on average (harder to remember nuance three sessions later).

## The risks, honestly stated

The same clinicians, asked about the downsides, name: **transcription errors** in clinical content (a mis-heard substance name, a missed negation), **loss of clinical voice** in notes (AI-drafted notes can sound homogenized), **over-reliance** (reviewing drafts with less rigor over time), and — critically — **privacy complexity** (audio leaves the room; where does it go, who has access, and how is it deleted?).

### The non-negotiables, before you adopt any clinical AI tool

1. A signed **Business Associate Agreement (BAA)** is in place before any PHI touches the tool. No BAA, no PHI. This is a HIPAA requirement, not a preference.
2. You have a **written informed consent process** that explains what the tool does, where data goes, how long it's retained, and the client's right to decline.
3. You have **read** (not skimmed) the vendor's data handling and model-training policy. Specifically: is your data used to train their models? Can you opt out? Can you delete on demand?

## WORKSHEET · 02

# Clinical AI vendor vetting

For each clinical AI tool you are considering, require answers in writing before signing. If a vendor won't answer one of these, that *is* your answer.

TOPIC	QUESTION	ANSWER / NOTES
<b>HIPAA / BAA</b>	Will you sign a Business Associate Agreement? Can I see the template?	
<b>Data location</b>	Where is PHI stored geographically? Is it encrypted at rest and in transit?	
<b>Model training</b>	Is my data used to train or improve your models by default? Can I opt out?	
<b>Subprocessors</b>	Which third-party AI providers (OpenAI, Anthropic, Google, AWS, etc.) do you use? Do they have BAAs with you?	
<b>Retention &amp; deletion</b>	How long is audio / transcript / draft stored? How do I delete on demand?	
<b>Accuracy &amp; error handling</b>	What's your documented word error rate? How do you handle hallucinations in drafts?	
<b>Breach history &amp; response</b>	Any reportable incidents in the last 3 years? What's your breach notification process?	
<b>Audit access</b>	Can I request a SOC 2 Type II report? HITRUST? Other independent certification?	
<b>Pricing &amp; lock-in</b>	Monthly cost, per-user cost, cost at scale? Export format if I leave?	



<b>Client-facing disclosures</b>	Do you provide consent language I can adapt? Is there a client-facing privacy page?	
----------------------------------	---	--

## PART 04

# AI in Business Practice

*Marketing, intake, admin — leverage and limits.*

## Where AI is genuinely useful on the business side

Unlike clinical AI, business AI has a much larger "obviously fine" zone — provided you keep a hard rule: **no PHI in non-BAA'd tools**. Everything else is a judgment call, and most of those calls come out in favor of using the tool.

### High-leverage, low-risk uses

**Marketing copy and content.** Website service pages, blog posts, newsletter drafts, social captions. Treat AI output as a first draft — your voice and clinical accuracy still have to be added by you.

**SEO research and structure.** Keyword brainstorming, FAQ drafting, meta descriptions, schema suggestions. Faster than any other method available to a solo practice.

**Policies and templates.** Informed consent drafts, cancellation policy rewrites, telehealth policies, supervision agreements — always reviewed by a qualified attorney or consultant before use.

**Admin drafting.** Responses to non-clinical emails, vendor inquiries, referral letters (with PHI removed), landlord correspondence, proposal drafts.

**Brainstorming and strategy.** Naming a group program, structuring a workshop, pricing analysis, competitive positioning.

### The line you shouldn't cross

Do not paste client-identifying details into consumer AI tools (the free or paid versions of ChatGPT, Claude, Gemini, etc., in their default consumer settings). Even if you think the tool "won't remember," you are outside HIPAA's protection without a BAA — and the risk is yours.

If you want to use a general-purpose AI on any content that includes PHI, either (a) use an enterprise / HIPAA-eligible tier with a signed BAA, or (b) rigorously de-identify the content before it leaves your system (remove names, dates, locations, ages, distinctive details).

**DO**

- Use AI to draft your 'About' page, then rewrite in your voice.
- Use AI to brainstorm blog topics for your specialty.
- Use AI to tighten a cancellation policy you already wrote.
- Use AI to generate keyword lists for SEO.
- Use AI to draft vendor-response emails.

**DON'T**

- Paste session notes into a consumer AI to "make it sound better."
- Use AI-generated clinical claims ("proven," "guaranteed," diagnoses).
- Publish AI text that makes outcome promises you can't substantiate.
- Let AI write your informed consent without legal review.
- Send an AI chatbot to handle crisis outreach.

**A quiet risk: the homogenization of your brand**

When every clinician's website uses AI to write 'I create a warm, collaborative, non-judgmental space,' potential clients stop being able to tell anyone apart. The clinicians who stand out in an AI-saturated field are the ones who sound *more* specific, more idiosyncratic, more human — not less. Use AI to get to a draft; use your own voice to get to a publish.

## PART 05

# Compliance, Ethics & Informed Consent

*The guardrails every AI-using clinician needs.*

## Four guardrails, in order of priority

You cannot wait for perfect regulatory clarity before acting — the regulatory picture is evolving in every state and at the federal level, and it will keep evolving for years. Instead, build your practice around four guardrails that are durable regardless of how specifics shift.

### Guardrail 1 - HIPAA and your Business Associate Agreements

Any vendor that creates, receives, maintains, or transmits PHI on your behalf is a business associate and needs a BAA. AI scribes, note-drafters that touch PHI, AI-enabled EHRs, AI-assisted billing services — all require BAAs. Keep a running list. Review annually.

### Guardrail 2 - State law and licensing board positions

Some state licensing boards have issued explicit guidance on AI use; many haven't yet. Check your board's website at least twice a year. Pay particular attention to: (a) documentation requirements, (b) advertising rules that AI-generated marketing might run afoul of, and (c) any supervision or disclosure requirements.

### Guardrail 3 - Professional ethics codes

The major professional bodies (ACA, APA, NASW, AAMFT, and others) have issued or are actively developing guidance on AI. The consistent themes across all of them: informed consent, client welfare as the primary duty, competence with the tools you use, protection of confidentiality, and honesty in marketing.

## Guardrail 4 · Informed consent that matches reality

If AI tools touch your work with a client in any way — a scribe in the room, an AI assistant drafting notes after the fact, a chatbot handling intake scheduling — that client deserves to know. Explicit, specific, written consent is the defensible standard. Generic "we may use technology" language is not.

### Minimum elements of an AI disclosure to clients

1. **What the tool does** (e.g., "records audio of sessions and produces a draft clinical note I review and edit").
2. **Where the data goes** (vendor name, general description of data handling).
3. **How long data is retained** and when it is deleted.
4. **Your right to decline** — without any effect on the quality or availability of your care.
5. **How to withdraw consent** later, and what happens to prior data if you do.
6. **A signature and date.** Verbal agreement alone is not enough.

WORKSHEET · 03

# Your compliance inventory

List every tool currently in your practice that could touch PHI or client-identifying information, and the status of each in HIPAA, consent, and documentation terms.

TOOL / VENDOR	PURPOSE	BAA?	CONSENT?	REVIEW DATE

Gaps I noticed (missing BAAs, missing consent, etc.):

---



---



---



---

First three gaps to close (with dates):

---



---



---



---

## PART 06

# Your AI Readiness Plan

*From reading to doing, in five steps.*

## A five-step plan you can run this quarter

Readiness is not a one-time event. It's a quarterly discipline. The plan below is deliberately modest: you are not overhauling your practice — you are moving it one honest step forward. Run it again next quarter, and the quarter after that.

### 01

#### Assess honestly

Complete Worksheets 01 (documentation), 03 (compliance), and the self-assessment below. Do not skip this step, even if you're tempted. You cannot improve what you haven't measured.

### 02

#### Pick exactly one pilot

One AI tool, one problem, one quarter. Most clinicians fail at AI adoption by trying to adopt three tools at once. The smallest successful pilot beats the most ambitious abandoned one.

### 03

#### Vet the vendor

Use Worksheet 02. Require written answers. If anything is unclear, stop — unclear is a decision in itself.

### 04

#### Run the pilot for 8–12 weeks

Set one or two specific metrics at the start (e.g., "minutes per note," "notes completed same-day," "client-reported satisfaction with consent process"). Log weekly; review at the midpoint.

### 05

#### Decide, document, disclose

At the end of the pilot: adopt, modify, or abandon. Write the decision and the reasons in your operations manual. Update client-facing disclosures if needed. Calendar the next quarterly review.

WORKSHEET · 04

# AI readiness self-assessment

Rate your practice honestly on a 1–5 scale (1 = not started, 5 = fully in place). Focus on the 1s and 2s first — they compound fastest.

DIMENSION	1	2	3	4	5	NOTES
Documentation habits meet payer-AI-ready standard						
Treatment plan goals are measurable & linked in notes						
CPT / time consistency across all claims						
BAA inventory is complete and current						
Written AI disclosure exists in client consent						
Clinical AI tools (if any) have been vetted in writing						
Business AI tools are used without PHI exposure						
Team / colleagues are trained on "no PHI in public AI"						
Licensing-board & ethics guidance checked within 6 months						
AI review is on the calendar quarterly						

**Lowest-scoring dimension I'll work on first:**

---



---

**One pilot I'll run this quarter (tool + problem + metric):**

---



---



---

WORKSHEET · 05

# Pilot plan (one page)

Fill this out **before** you touch the tool. A pilot without a plan becomes a habit without evidence.

**Tool & vendor**

---

---

**Problem I'm trying to solve**

---

---

**Specific success metric(s) — numeric if possible**

---

---

**BAA signed? Consent updated? Disclosure drafted?**

---

---

**Clients included (all? new only? those who consent?)**

---

---

**Start date / Midpoint review / End date**

---

---

**What would make me stop the pilot early?**

---

---

**Who else needs to know, and how are they being informed?**

---

---

## PART 07

# The Readiness Checklist

60+ items, organized by domain. Revisit quarterly.

This is the practical heart of the workbook. It's designed to be photocopied, printed, or revisited digitally each quarter. A checked box next to every item isn't the goal; **a thoughtful pass through every item** is.

## Documentation & payer readiness

- Every progress note names a specific intervention (not just "supportive" or "processed").
- Every progress note describes client response to the intervention.
- Every progress note includes at least one observable symptom with severity or frequency.
- Every note references at least one current treatment plan goal.
- CPT code billed matches documented session time in every claim.
- Treatment plans have measurable, observable, dated goals.
- Treatment plans are reviewed and updated at least every 90 days (or per payer rules).
- Diagnosis is re-justified in notes whenever it changes or every 6 months at minimum.
- I have read each of my primary payers' current medical-necessity criteria.
- I can produce a complete, coherent chart for any active client within 24 hours.

## HIPAA & privacy foundation

- A signed Business Associate Agreement (BAA) is on file for every vendor that touches PHI.
- My BAA list is reviewed at least annually and dated.
- I have written policies on acceptable use of AI tools in my practice.
- No PHI is ever pasted into a consumer AI tool without a BAA.
- If I use enterprise AI for PHI, I have confirmed in writing the BAA covers that usage.
- I understand each AI vendor's data retention and deletion policy.
- I know whether each AI vendor uses my data to train models, and I have opted out where possible.
- Two-factor authentication is enabled on every AI tool with PHI access.
- My EHR's AI features (if any) have been reviewed, not just auto-enabled.
- A breach response plan exists and names a responsible person.

## Clinical AI tool evaluation

- Every clinical AI tool in use has gone through Worksheet 02 vendor vetting.
- I personally reviewed (not skimmed) each vendor's privacy and data-training policy.
- I know the approximate word error rate / accuracy of any AI scribe I use.
- I review every AI-drafted note before signing, without exception.
- I have caught and corrected at least one AI error and documented the process.
- I have a plan for what to do if the AI tool goes down during or after a session.
- I do not allow AI tools to auto-sign notes on my behalf.
- Clinical AI outputs are treated as drafts, not final documentation.
- My supervisor / consultation group knows which AI tools I use.
- I re-evaluate each clinical AI tool at least every 6 months.

### Informed consent & client disclosure

---

- My informed consent includes specific, named AI tools — not just "technology."
- Consent explains what each tool does, where data goes, and how long it's kept.
- Consent explains the client's right to decline without any effect on their care.
- Consent explains how a client can withdraw AI consent later.
- Consent is in writing and signed before any AI tool touches the client's session.
- I verbally walk new clients through the AI portion of consent, not just hand them a form.
- Existing clients were re-consented before I added an AI tool to their care.
- My intake process includes a clear "I decline AI tools" option.
- I document the consent conversation in the client's chart.
- I have a script for clients who ask, "Are you using AI with me?"

### Business & marketing AI

---

- AI-drafted website and marketing copy is reviewed by me before publishing.
- No outcome guarantees, "proven," "guaranteed," or superlative claims appear in AI-assisted copy.
- Credentials, licenses, and modalities listed on AI-assisted pages are verified accurate.
- My voice and specificity — not generic AI phrasing — are visible on my site.
- AI-drafted client emails do not contain PHI pulled from charts.
- AI-generated blog content is fact-checked against current clinical literature where relevant.
- I do not use AI-generated images to misrepresent my practice (fake offices, fake teams, etc.).
- My staff and contractors are trained on "no PHI in public AI."
- I have a written policy for any contractor who uses AI on my behalf.

- Invoices, marketing, and admin tools have been reviewed for AI data sharing I didn't intend.

### Ethics, licensing & professional standing

---

- I have read my licensing board's current guidance on AI (or confirmed there is none).
- I have read my professional association's AI position (ACA, APA, NASW, AAMFT, etc.).
- I have completed at least one CE on AI in mental health in the last 12 months.
- AI use is on the agenda for my consultation group at least once a year.
- I do not represent AI-assisted work as fully human-produced in marketing.
- I do not use AI to practice outside my scope (diagnosis, populations, modalities).
- I know whether my malpractice carrier has a position or requirement on AI use.

### Quarterly review discipline

---

- AI review is a recurring calendar event — not something I do when I remember.
- Each quarter, I re-run Worksheet 04 (self-assessment) and log the score.
- Each quarter, I check each payer's medical-necessity criteria for updates.
- Each quarter, I check each AI vendor's data policy for changes.
- Each quarter, I confirm every BAA on file is still current and signed.
- Each quarter, I review at least 5 random recent notes against my documentation standard.
- Each year, I ask one peer or consultant to review my AI policies and consent.

---

*Total checklist items: 64. Don't aim for all checked in one quarter. Aim for honest progress and a dated review.*

**CLOSING**

# Competence, not certainty.

If you worked through even half of this workbook, you are already ahead of most of your field on this question. Not because you have all the answers — you don't, and nobody does yet — but because you have a **structured, honest way to ask them.**

The counselors and group practices who will thrive in the AI era won't be the ones who adopted earliest or who resisted longest. They'll be the ones whose documentation is clean enough to survive algorithmic review, whose tools are chosen deliberately, whose clients know exactly what's happening with their information, and whose practice reviews this every quarter — quietly, un sentimentally, and without drama.

Put this workbook somewhere you'll see it in 90 days. Open it. Run through the checklist again. Update your scores. Close one more gap.

That's the work. That's the whole thing.

**A final, important note**

This workbook is educational. It is not legal advice, compliance advice, clinical advice, or a substitute for consultation with a qualified attorney, supervisor, ethics board, or professional association. Laws, regulations, and professional standards regarding AI are evolving rapidly — confirm any specifics against current authority before relying on them in your practice.

*Thank you for the work you do.*

# You don't need to predict the future. You need to prepare for it.

AI is already reshaping how payers review claims, how clinicians document sessions, and how practices market their services.

This workbook gives you a structured, unhurried path through:

- What your payers are doing with AI — and what it means for your documentation and medical necessity language.
- How to evaluate clinical AI tools (scribes, note-drafters, decision support) against HIPAA, ethics, and clinical fit.
- How to use AI in the business side of practice without putting PHI, your license, or your brand at risk.
- A step-by-step readiness plan, worksheets, and a comprehensive checklist you can return to each quarter.

## INCLUDES

- 5 fill-in worksheets
- Vendor vetting template
- 60+ item readiness checklist