

Is Your AI Ready to Survive an Audit?

The 5-Minute AI-Readiness Check + Worksheet

Your team is already using AI. The question is whether you could prove what it did if a client, a regulator, or opposing counsel asked. Answer each question below — score 1 point for every “No / Not sure.”

PART 1

The Check

1. The perimeter. Can you name exactly where your data goes when your team uses an AI tool?	☐ Yes	☐ No / Not sure
2. Shadow AI. Could a staff member paste a confidential document into a free AI tool today without anyone knowing?	☐ No, covered	☐ Yes / Not sure
3. Vendor agreement. For any AI tool touching sensitive data, has the vendor signed a BAA / DPA?	☐ Yes	☐ No / Not sure
4. Audit trail. Could you produce a log of every question your AI was asked about a matter and what it produced?	☐ Yes	☐ No / Not sure
5. Grounding. Does your AI answer from your real documents — or sometimes make things up?	☐ Grounded	☐ Makes things up / Not sure
6. Human in the loop. Is there a human reviewing AI output before it reaches a client or drives a decision?	☐ Always	☐ Not consistently
7. The policy. Do you have a written AI-use policy your team knows and follows?	☐ Yes	☐ No / Not sure
8. The rulebook. Do you know which rules govern your AI use (Rule 1.6, HIPAA, GDPR, FINRA)?	☐ Yes	☐ No / Not sure
9. Access control. Is access to AI systems and the data they reach limited by role?	☐ Yes	☐ No / Not sure
10. Ownership of output. Do you know who owns, and who can see, what your AI produces?	☐ Yes	☐ No / Not sure

YOUR SCORE	Tally every “No / Not sure”	
0–2	On Track	Ahead of most. The risk is the exposure you can’t see — a short review confirms it and catches blind spots.
3–6	At Risk	Real, specific exposure. One pasted document or client questionnaire from a problem. The most common, most fixable range.
7–10	Exposed	Likely already out of compliance somewhere, and you couldn’t prove otherwise. Put a defensible version in place fast.

Whatever your score, the next step is the same: knowing you have gaps and knowing which 2–3 to fix first are very different. That’s the **AI Readiness Session**.

PART 2

The Worksheet

Fill this in and you've done the first hour of the work — it's exactly what we build out together in an **AI Readiness Session**.

A. Tool inventory

Every AI tool your team uses — or might be using.

TOOL	WHO USES IT	DATA IT COULD TOUCH	VENDOR AGREEMENT?

B. Sensitive-data map

What protected / regulated information exists? (client files, patient records, financials, privileged material, PII, IP)

C. The 5-component gap check

Rate each pillar: **R** = missing **Y** = partial **G** = solid

COMPONENT	WHAT IT MEANS	RATING
Data perimeter	Sensitive data stays inside your control	R ☐ Y ☐ G ☐
Grounded retrieval	Answers from your real documents, not guesses	R ☐ Y ☐ G ☐
Audit trail	Every query, source, and output is logged	R ☐ Y ☐ G ☐
Human in the loop	A person reviews consequential output	R ☐ Y ☐ G ☐
Governance	Policy, access control, and ownership are defined	R ☐ Y ☐ G ☐

D. Your top 3 exposures

From the above, the three things that worry you most.

1.

2.

3.

E. Where AI would actually help

The two workflows where AI would save the most time or money if it were safe to use.

1.

2.

Now make it a plan.

Bring this worksheet to an **AI Readiness Session** — I read it first, then we spend one focused hour on your biggest exposures and the two or three things to fix first. **\$250**, a complete session, no pitch. Ready to go deeper? The **AI Readiness Audit** is a 1–2 week engagement, credited toward your build.

Book your AI Readiness

Session: calendly.com/kylecedrichanna/ai-readiness-intro · kyle@kylecedrichanna.com