THE AI ACCOUNTANT
Your Practice, powered by AI

FREE RESOURCE

The AI Disclosure Position Kit

From “no position” to “written position” in one sitting: the scope test, the exposure matrix, the supplier questions, and the statement to put your name on.

Peter McCarroll, CPA · The AI Accountant
theaiaccountant.ai · August 2026

HOW TO USE THIS KIT

One sitting. One written position.

In four weeks in mid-2026, AI disclosure went from an ethics debate to infrastructure. Substack now scans posts and shows readers a percentage estimate of how much a human wrote. LinkedIn added a report option for AI slop. The EU made disclosure a legal duty. And Anthropic put a watermark in every word Claude writes, worldwide, with the other major model developers following. Nobody in the profession voted on any of it.

Here is the good news: your firm almost certainly faces no legal risk, and the position that protects you is one you can write today. Work through the sections in order. By the end you will know which role your firm occupies, which outputs are even in scope, where your detection exposure actually sits, and you will have a disclosure position statement with a name on it.

What's inside

SECTION 1

The argument on one page

Detection is infrastructure now.

The scan is not coming; it shipped. Substack's reader-facing AI detection covers posts published on or after July 21, 2026. LinkedIn's "seems like AI slop" report option arrived in late July, and flagged posts are suppressed beyond the poster's own network. The EU AI Act's transparency obligations took effect August 2. Anthropic's watermark shipped worldwide on August 14: a statistical fingerprint hidden in the pattern of word choices, invisible to any reader, detectable by software built to look for it. Assume every piece of text your firm produces can be scanned, and that the scan is free.

Provenance is not responsibility.

Four instruments now point at your writing, and they measure four different things. Style classifiers guess from prose patterns, and they are known to flag genuinely human writing. Watermarks record provenance: which words a model produced. The EU rule tests responsibility: did a human review it, and is a named person accountable? And "slop" measures care: was this worth anyone's attention. The public is about to read all four as one dial. They are not one dial, and the needle points backwards. The watermark weakens on short, factual, and edited text, so hand-paraphrased junk sheds its flag and passes clean, while long-form content produced under the heaviest editorial control flags as maximally AI. The instrument is strongest exactly where the human was most involved.

The position is the lever.

You cannot remove the flag without burning the hours AI handed back; retyping and paraphrasing is work that adds no value, performed to hide the tool. What you control is the story that arrives with the flag. The EU has already defined the safe harbor: human review plus named editorial responsibility. Substack has already built the slot: the "How I make this" statement shown whenever a reader scans your work. What is missing at most firms is the sentence to put in it. This kit gets that sentence written, with your name on it, today.

> **The one-line summary**
>
> The firms that get hurt over the next two years won't be the ones whose content flags as AI. They'll be the ones whose flag arrives with no story attached, or with a story the flag contradicts. "100% handcrafted" is now a checkable claim.

SECTION 2

The four roles

The EU AI Act's disclosure rules don't attach to “companies that use AI.” They attach to specific roles, and they attach per AI system, not per firm. Find your row before you spend a dollar on compliance.

Role	Who's in it	Article 50 duty	What it means for you
1. Provider: model developer	Anthropic, OpenAI, Google, and the other model developers.	Mark AI output in machine-readable form (Art. 50(2)). The watermark is this duty, discharged.	Nothing to do. This row is why the watermark exists.
2. Provider: AI-embedded software	Any product that puts an AI system on the market under its own name or trademark, whether the model inside is its own or rented.	The same marking duty, for that feature.	Your ledger, workflow, or tax software carries this duty for its AI features. Their problem; your renewal-time question (Section 6).
3. Deployer: firm using AI	An organization using an AI system under its own authority (Art. 3(4)).	No marking duty. Only Art. 50(4) can ever reach you, and only for published content.	Almost certainly your row, for almost everything your practice does. Run Section 3.
4. Deployer that publishes AI output	Firms whose AI output is itself the product: benchmark data, industry reports, public tools.	Still a deployer, but the published-content limb of Art. 50(4) is live.	Closest to the line. Run Section 3 on every published output.

The capability test for the provider line

Selling a subscription to AI-generated reports keeps you a deployer. Putting an AI-powered tool in customers' hands under your brand (a business-plan generator running on a rented model, say) makes you the provider of that system, marking duty included. The line is crossed on capability, not content, and “free of charge” doesn't change the analysis.

The per-activity rule

Roles attach to each AI system separately. Becoming the provider of one product doesn't reclassify anything else your firm does. Classify system by system.

Where typical firm activities land

Using Claude or ChatGPT to draft client emails, workpapers, or memos	Deployer.

Running the AI features inside your ledger, workflow, or tax software	Deployer; the software company is the provider of that feature.
Publishing an AI-drafted blog, newsletter, or LinkedIn post	Deployer; the Section 3 scope test is yours to run.
Selling a subscription to AI-generated reports or benchmarks	Still a deployer, with the published limb live.
Launching a tool under your brand where customers generate output themselves	Provider of that tool, with its marking duty. Still a deployer everywhere else.
Offering a free assessment or calculator that emails back a personalized AI report	Provider if the model writes the prose the reader receives; deployer if it only scores and routes among human-approved blocks. Either way the report isn't published, so no disclosure duty attaches to it.

If any provider row is yours, that is a product-compliance project that belongs with legal counsel, not this kit. Everything from here on assumes the role your practice is almost certainly in: deployer.

SECTION 3

The five-gate scope test

Run each output type through the gates in order. The first "no" ends the analysis: out of scope. Under Article 50(4), disclosure is owed only when every gate is passed and no exemption applies.

Gate 1 — Role

For this output, is your firm a deployer? If one of Section 2's provider rows applies, stop here: that is a different set of duties, and it belongs with counsel.

Gate 2 — Published?

Is the content published? Client deliverables exit here: the monthly pack, the tax return, the advisory memo, the reminder email. Not published means not in scope, no matter how much AI touched it. That is not an interpretation; it is what the text says.

Gate 3 — Public-interest purpose?

Is it published "with the purpose of informing the public on matters of public interest"? GRAY ZONE: the term is undefined and genuinely unsettled. Pure marketing arguably is not; commentary on tax law changes arguably is. Record your honest answer, and don't try to resolve the ambiguity from your desk.

Gate 4 — Used in the Union?

Is the output "used in the Union" (Art. 2(1)(c))? GRAY ZONE: whether that requires targeting EU readers or merely being readable there is unsettled. EU-based firms answer yes and move on.

Gate 5 — The editorial-control carve-out

Has the content "undergone a process of human review or editorial control," with a natural or legal person holding editorial responsibility for the publication? If yes: exempt. This is the exit written for professional firms, and it is the most consequential sentence in the regime.

Caught by all five gates?

Two ways to comply. Path 1: say, clearly and visibly with the content, that it is AI-generated. No filing, no regulator to notify, no prescribed wording; a label. Path 2: restructure the workflow so the carve-out applies, with real human review and a named person holding editorial responsibility. Take Path 2. It is also the answer to the watermark, the platforms, and your clients.

The worksheet

One row per output type. Write yes or no in each gate column; the first no is your verdict.

Output type	G1 role	G2 publ.	G3 pub-int.	G4 EU	G5 carve-out	Verdict
Example: firm blog post	Deployer	Yes	Unsettled; assume yes	Unsettled; assume yes	Yes: review + named person	**EXEMPT**

The penalty, for perspective. The number in the Act is EUR 15,000,000 or 3% of worldwide annual turnover (Art. 99(4)(g)), with reduced caps for SMEs (Art. 99(6)). It is a number built for model developers, and it sits behind every gate above.

SECTION 4

The two-axis exposure matrix

The law is one instrument; the watermark and the platforms are others, and they do not measure responsibility. Place each output on two axes: who produced the words, and who stood behind them before they moved. Cross the result with private versus published and you know both your legal position and your detection exposure.

	Auto-sent / auto-published (no instance review)	Reviewed before it moved
AI produced the words	**Highest exposure.** Private: no EU scope, but provenance marks can surface in client-facing text. Published: reaches the disclosure test with no carve-out. Fix the workflow or label it.	**The professional quadrant.** Private: out of scope, strong position. Published: exempt via the carve-out, and still maximally flagged by the watermark. Attach your story (Section 5).
A human produced the words	Template letters and mail-merge. No AI issue; ordinary quality risk only.	Traditional work. Note: style classifiers still flag some genuinely human writing, and a false flag with no position is still a bad meeting.

Common outputs, pre-placed

Auto-generated client emails and reminders	AI × auto, private. Out of legal scope; watch for marks in client-facing text.
Reviewed monthly summaries and advisory memos	AI × reviewed, private. Out of scope; your strongest quadrant.
Auto-published social content	AI × auto, published. Highest exposure on the board: reaches the test with no carve-out. Restructure or label.
Governed long-form articles and newsletters	AI × reviewed, published. Exempt under the carve-out and maximally flagged by the watermark. A written position is mandatory here.
Automated personalized deliverables (assessments, calculators, proposal generators)	System-level accountability with no instance review. Put a by-design statement in the deliverable itself: “This report is generated automatically; the methodology was designed and reviewed by [Name].”

The deciding test for tool builders

Does the model choose the words the reader receives, or does it only score and route among blocks a human wrote and approved in advance? The first is provider territory (Section 2). The second stays deployer.

SECTION 5

The template position statement

> *"We use AI in production. Every deliverable passes human review before it reaches you. [Name] holds responsibility for what we publish and deliver."*

How to adapt it

Swap [Name] for the partner or owner who actually signs. Make sure the middle sentence is true before you publish it: if some deliverables don't pass human review, either change the workflow or scope the sentence honestly ("every deliverable we publish," "every client-facing report"). Keep it to three sentences; a position that runs to a page reads as a defense.

Where to put it

Substack's "How I make this" field, verbatim: it appears whenever a reader runs a scan. Your website footer or about page. The technology clause of your engagement letter. Proposals and onboarding decks. One position, used everywhere; a position that varies by channel reads as improvisation.

The variant for automated deliverables

For assessments, calculators, and proposal generators, where accountability is at the system level rather than per instance:

> *"This report is generated automatically. The methodology was designed and reviewed by [Name], who holds responsibility for it."*

The test of a good position

A client who runs a scan, sees a 100% AI flag on your article, and then reads your statement should find agreement, not contradiction. The flag says a model produced the words. Your statement says who controlled them and who answers for them. Both can be true, and together they are the whole story.

SECTION 6

Questions for your software suppliers

Section 2's second row is your suppliers' duty, and their answers land in your client-facing output. Ask these at renewal, not after a client asks you.

1 **Which features of your product generate text or other content with an AI model, and whose model is it?**
Notes:

2 **Do you mark or watermark AI-generated output in machine-readable form? Which method?**
Notes:

3 **What is your Article 50 position? Are you a signatory to the EU Code of Practice on AI content transparency?**
Notes:

4 **Will marks appear in output we deliver to clients (reports, letters, summaries)? Where, exactly?**
Notes:

5 **If you change the underlying model, does marking behavior change with it? How will you notify us?**
Notes:

6 **For systems on the market before August 2, 2026: what is your compliance timeline against the December 2, 2026 transition deadline (Art. 111(4))?**
Notes:

7 **Who at your company answers AI compliance questions, by name?**
Notes:

SECTION 7

Disclosure do's and don'ts

Under-disclosure and over-disclosure both cost you, in different currencies. The line between them is the review layer.

DO	DON'T
Write one position and use it everywhere. A position that varies by channel reads as improvisation.	**Don't blanket-label everything "AI-generated."** Over-disclosure signals the absence of the review layer, and it forfeits the carve-out's story while gaining nothing.
Name a real person. The EU carve-out, your clients, and your own team all need the same thing: someone identifiable standing behind the work.	**Don't claim handcraft you can't prove.** Undisclosed AI plus a "handcrafted" claim is the one position measurably worse than no position.
Keep the review real. The carve-out requires a process of human review or editorial control, not a sentence claiming one. If the sentence is ahead of the workflow, fix the workflow.	**Don't hand-paraphrase to strip watermarks.** That is work that adds no value, performed to hide the tool, and it burns the hours AI handed back.
Audit your old claims. Find and retire any "written by humans" or "100% handcrafted" language you can no longer prove. It is now a checkable claim.	**Don't toggle detection off where readers can see the switch.** "AI detection unavailable" is its own kind of disclosure.
Ask suppliers before renewal. Section 6 takes fifteen minutes at renewal and one awkward month after a client scan.	**Don't wait for the first client scan.** The story you tell after a flag reads as an excuse. The same story told before reads as a policy.

SECTION 8

Citations and disclaimer

EU AI Act, Regulation (EU) 2024/1689, Art. 2(1)(c)	Reach: the Act covers third-country providers and deployers “where the output produced by the AI system is used in the Union.”
Art. 3(3) — “provider”	Develops an AI system, or has one developed, and places it on the market “under its own name or trademark, whether for payment or free of charge.”
Art. 3(4) — “deployer”	A person or body “using an AI system under its authority,” outside personal non-professional activity.
Art. 50(2)	Provider duty to mark synthetic output in machine-readable form. The watermark implements this.
Art. 50(4), text limb	Disclosure for published AI-generated text “with the purpose of informing the public on matters of public interest,” except where the content “has undergone a process of human review or editorial control and where a natural or legal person holds editorial responsibility for the publication.”
Art. 99(4)(g), 99(6)	Administrative fines up to EUR 15,000,000 or 3% of worldwide annual turnover; reduced caps for SMEs.
Art. 111(4)	Transition for systems already on the market before August 2, 2026 (providers): compliance by December 2, 2026.
Anthropic, “How Claude's text watermark works” (Aug 14, 2026)	anthropic.com/news/claude-text-watermark. Worldwide application (“we don't yet have a durable way to scope it by region”); roughly 190 signatories to the EU Code of Practice on AI content transparency; watermark weakens on short, factual, and edited text; detection API planned; method based on Google DeepMind's SynthID-Text (Nature, 2024).
Substack Help Center, “How can I detect AI on Substack?” (updated Aug 9, 2026)	Reader-facing detection for posts published on or after July 21, 2026, using Pangram; percentage estimate; creator “How I make this” statement shown on scan; per-post disable shows readers “AI detection unavailable.”
LinkedIn “seems like AI slop” report option (late July 2026)	Per secondary reporting (MediaPost, Jul 31, 2026; TechCrunch; Forbes, Aug 10, 2026): report option added; flagged posts suppressed beyond the poster's network. No primary LinkedIn announcement as of Aug 15, 2026.

Key dates

July 21, 2026: Substack reader-facing detection (posts on or after this date). Late July 2026: LinkedIn slop reporting. August 2, 2026: EU AI Act transparency obligations apply. August 14, 2026: Anthropic watermark ships worldwide. December 2, 2026: transition deadline for pre-existing systems (providers).

Disclaimer

This kit is analysis and education, not legal advice. Article 50's scope turns on facts specific to your firm, and two terms in this analysis (“used in the Union” and “matters of public interest”) are genuinely unsettled. If real money or regulatory exposure rides on the answer, take the question to counsel.