
AI in Casework

Where artificial intelligence touches our work and where it does not, what happens to prompts and inputs when a matter concludes, and why no AI output ever becomes part of a case file.

Effective August 18, 2026 · Next scheduled review February 18, 2027

Institutions are beginning to ask vendors where AI sits in their process, and they are right to. An investigative record is only as defensible as the reasoning behind it, and reasoning a practitioner cannot account for is not reasoning at all.

So we state our position plainly rather than waiting to be asked.

The rule that governs everything else

A human practitioner reaches every finding, writes every determination, and is accountable for both. No AI system decides credibility, weighs evidence, reaches a conclusion, or drafts a finding that goes out under a practitioner's name.

That is not a limitation we expect to relax as the technology improves. The reason a finding holds up is that a named person with relevant expertise made it and can explain it under questioning, and no model can be cross-examined.

Nothing AI produces enters a case file

AI output is never part of the record. It is not attached, quoted, summarized into, or cited in a report, a determination, a notice, or an appeal record. Where a practitioner uses a tool to organize their own thinking, what reaches the file is the practitioner's own work, verified against the underlying evidence.

This is deliberate and it matters for a specific reason. A file that contains AI output invites a challenge about how that output was produced, on what data, and with what error rate. That is a fight about the tool rather than about the facts, and it is not a fight we will hand to a respondent's counsel.

Who processes it, and under what terms

We hold an enterprise agreement with Anthropic and use Claude under its commercial terms. That matters for three reasons. Anthropic acts as a processor rather than a controller, so the information stays ours. Inputs and outputs are excluded from model training by contract rather than by a setting somebody could switch. And the arrangement is an enterprise one, not a consumer account, which is where most organisational exposure actually comes from.

We do not use consumer AI accounts for any work connected to an engagement, and practitioners are not permitted to. That is a term of engagement rather than a guideline.

Prompts and inputs are deleted when the matter concludes

Prompts, inputs, intermediate outputs, and any working context are held only while the matter is live. We keep our own working copies on encrypted systems under our control, and our provider retention is configured to the shortest term available to us.

When a matter concludes, they are deleted. They are not retained under our records schedule, because they are not records of the matter. They are working scaffolding, and scaffolding comes down when the building is finished.

The consequence is worth stating: if you ask us later what was in a prompt on a concluded matter, we will not be able to tell you, because it no longer exists. The file will contain the practitioner's work and the evidence it rests on, which is what a reviewer needs.

What is retained, and what is not

These are two different categories and the difference is deliberate.

- **Retained.** Investigation reports, determinations, notices, and evidence exhibits are records of the matter. They are held under our records retention schedule, ordinarily for seven years, because a reviewer may need to examine them years later.
- **Not retained.** Prompts, model inputs and outputs, and working context are not records of the matter. They are deleted when it concludes and appear in no schedule.

The line is drawn on a single question: does a reviewer examining this matter need it to understand what the institution did and why? A report and an exhibit answer that question. A prompt does not, because nothing in the file rests on one.

What we never put into a model

- Party or witness names, or details that identify a person by inference.
- Evidence, statements, interview recordings, or transcripts.
- Draft findings, determinations, or deliberative material.
- Institutional records provided to us under an engagement.
- Anything covered by privilege where we are engaged through counsel.

Where a tool would genuinely help with a task that touches this material, the answer is that the practitioner does the task.

Where AI is actually useful to us

We are not against the technology. Used honestly, it helps with work that does not touch a case record:

- Research on published regulations, agency guidance, and case law, verified against the source before it informs anything.
- Drafting and editing our own published material, this website, and internal documents that contain no client information.
- Structuring training content and curriculum, reviewed by the practitioner who will teach it.
- Administrative work: scheduling, formatting, summarising our own public writing.

No training on your information

Under the commercial terms we operate on, our inputs and outputs are not used to train models. That exclusion is contractual and it is not subject to a toggle in a settings page.

Two things are worth saying honestly about that. A contract term binds a vendor legally, and it is not something either of us can verify technically from the outside. And safety systems at any major provider may retain flagged content for a period regardless of the arrangement. We will show you the terms we operate under if your review needs them.

OnePointe

OnePointe is a separate company. We support and deploy it as our recommended case management platform, and we do not own it. Where an institution uses OnePointe, the platform's own terms, security documentation, and AI disclosures govern that relationship directly, and you should read them rather than relying on ours.

This policy governs what our practitioners do, in whatever system they are working in.

Disclosure

If you want to know whether AI touched any part of an engagement, ask and we will tell you, specifically. We will also tell you before the engagement begins if we anticipate using it for anything at all.

Where an institution prefers that no AI tool be used at any point in its engagement, we will work that way and say so in the engagement letter. Nobody has to justify that preference to us.

Bias

Civil rights work is the wrong place to be casual about automated bias. A system trained on historical decisions will reproduce the patterns in them, and the patterns in institutional discipline records are exactly the patterns we are engaged to examine.

That is a substantial part of why findings stay with people. It is also why we will not recommend any tool that scores, ranks, or predicts outcomes for parties in a civil rights matter.

If we get this wrong

If a practitioner uses a tool outside this policy, we will tell the institution, assess whether any part of the record is affected, and redo the work at our cost where it is. We would rather report it than have it surface in a challenge.

How to reach us

Fractional Coordinator, Inc.

830 NE Holladay St., Portland, OR 97232, United States

Office: (503) 395-0687 · Fax: (503) 395-0815

Questions about this policy: hello@fractionalcoordinator.com