

Dear Partner, We Need to Talk About AI

March 23, 2026 | 9 min read



An honest conversation about what AI can and can't do for your law firm — and what you should actually be worried about.

Look, I'm not going to dress this up in consultant-speak. I build software for law firms. I sit in your offices. I see what's actually happening on the ground. And there are things I need to tell you that nobody else will — because the AI vendors want your money and the tech press wants your clicks.

So let's have a real conversation. Imagine we're at Java House, it's 3pm, and you've just asked me: "Lee, what should I actually be doing about this AI thing?"

Here's what I'd say.

First — Your Associates Are Already Using AI. You Just Don't Know About It.

This is the conversation I have with every managing partner, and it always goes the same way. I ask: "Do your associates use ChatGPT?" The partner says no, or says they're not sure. Then I ask the associates privately. Every. Single. One. Is using it.

They're copying client instructions into it. They're pasting contract clauses. Settlement figures. Witness statements. Sometimes entire briefs. And they're doing it on the free version — the one where the terms of service basically say "***we can use your inputs to train our models.***"

Do you understand what that means? Your client's privileged information — their M&A terms, their exposure in a dispute, their financial records — is sitting in a training dataset. Today it's invisible. Tomorrow, some version of it could surface when someone else asks the right question. That's not science fiction. That's how these models work.

Your biggest AI risk right now isn't that robots will replace your lawyers. It's that your lawyers are feeding client secrets into tools you don't control, don't monitor, and probably don't even know about.

And here's the thing — you can't just ban it. I've seen firms try. They send a memo saying "no AI tools." You know what happens? People keep using them. They just stop telling you. A ban doesn't stop usage. It stops *disclosure* of usage. Which is worse.

So What Do You Actually Do?

Three things. Today. Not next quarter.

One — accept reality. Your people are using AI whether you like it or not. The horse has left the stable. Your job isn't to chase the horse. Your job is to build a better stable.

Two — give them a sanctioned tool. Get a Claude or ChatGPT enterprise subscription. Something where inputs aren't used for training, where you control data retention, where there's an audit trail. The cost is maybe fifty, sixty thousand shillings a month for a team plan. Compare that to one confidentiality breach landing before the Advocates Complaints Commission. That's not even a hard calculation.

Three — write a one-page AI policy. Not a 40-page compliance manual nobody reads. One page. What goes in, what never goes in (client names, case numbers, privileged communications without anonymisation), and what verification is required before anything AI-touched leaves the firm. Pin it next to every screen.

Is AI Just... Chatbots? Is That All There Is?

No. And this is where most lawyers get confused, because all the noise is about chatbots. ChatGPT this, Gemini that. So everyone thinks AI means "a chatbot I type questions into."

Let me be clear: **a chatbot by itself is almost useless to a law firm.** A chatbot is just an interface. It's like having a really smart person locked in a room with no files, no systems, no access to your client records, no connection to your calendar, no way to generate your invoices. They can talk. That's it.

What you actually need is **AI connected to a system.** An AI agent that can pull up a client file, draft a billing narrative, flag a missed deadline, organise your matter records — that's useful. But it needs a platform underneath it. It needs your practice management system, your document management, your billing engine. Without that, you're just having a conversation with a very articulate parrot.

So when someone tells you "just get ChatGPT and you're sorted," they're selling you half the story. The AI is the brain. The system is the body. You need both.

What AI Is Actually Good At: Drafting

I'll give credit where it's due. For drafting — and I mean first-draft generation — AI is genuinely impressive. Demand letters, standard contract clauses, board resolutions, statutory notices, client correspondence. The formulaic stuff. AI can cut your drafting time by 70–80% on routine documents.

But — and this is a big but — it's a *first draft*. Every single time. You don't paste the output into a court filing. You don't email it to the client without reading every line. As I wrote in [my earlier piece](#): **you cannot outsource your license to an algorithm.**

The right way to use AI for drafting isn't to ask it to "write me a demand letter." That gives you generic rubbish. The right way is to build template systems — your firm's best templates, your standard clauses, your proven structures — and use AI to populate them with case-specific variables. You review, you add the human judgment, you sign. Ninety percent time saved, zero percent hallucination risk.

Now the Honest Part: AI Is Not Good at Legal Research. Not Yet.

This is where I'm going to say something unpopular, and I need you to hear it.

If you're using AI to do legal research in Kenya right now, you are probably slower than if you did it the old-fashioned way. And you might be wrong on top of it.

Here's why. Kenya's case law isn't deeply embedded in these AI models. The Kenya Law Reports database is growing, yes, but it's not what these models were trained on. So when you ask an AI to find Kenyan precedent on, say, the enforceability of penalty clauses under the Law of Contract Act, it might give you something that sounds perfect — a realistic case name, a convincing citation, a plausible ratio — and the case simply does not exist. It made it up. The technical term is "hallucination." The practical term is "career-ending if you cite it in court."

The other problem is currency. New rulings come out from the Court of Appeal, the Supreme Court, the Employment Court, all the time. AI models have knowledge cutoffs. Even with web search, the tool might miss a critical 2025 ruling that completely changes the landscape of your dispute. You cannot afford to not know about that ruling because your AI didn't find it.

My honest advice to any advocate: do your research the traditional way. Browse the books. Pull actual reports from Kenya Law. Read the judgments yourself. Build your opinion on verified sources. Then — and only then — bring AI into the picture.

How to Actually Use AI for Research (Without Getting Burned)

Here's the workflow I recommend, and I genuinely think this is the safest, most effective approach available today:

Step 1: Do the research yourself. Go to Kenya Law. Pull the case reports. Check the physical volumes if you need to. Read the statutes. Build your legal opinion independently of AI. This is the work. There's no shortcut.

Step 2: Draft your opinion. Write up your analysis, your arguments, your recommended position. Do this with your own legal mind.

Step 3: Now bring in AI — but as your moot court opponent, not your research assistant. Pass your drafted opinion through AI and ask it: "What arguments would opposing counsel make against this?" or "What angles haven't I considered?" or "Where is this argument weakest?" You're not asking it to find sources. You're asking it to stress-test your thinking.

Step 4: Use NotebookLM. This is my specific recommendation. Upload all the actual sources you used — the judgments, the statutes, the commentary — into Google's NotebookLM. Then use it to analyse that closed universe of documents. Ask it to spot issues you might have missed, contradictions between your sources, or arguments you haven't considered. Because it's working only with documents you've uploaded and verified, it can't hallucinate a source that doesn't exist. That's the key insight most people miss.

Think of AI in research not as your research assistant, but as the advocate on the other side of the moot. It pokes holes in your argument after you've already built it. That's where it adds value. That's where it's safe.

Beyond Drafting and Research — What Else Can AI Actually Do for You?

This is the question every advocate asks me, and it's the right question. Because if AI is only about drafting and research, it's a limited tool. But here's where it gets interesting — **and here's where you need systems, not just chatbots.**

Client Intake and Triage

How much of your time goes to initial consultations where the client talks for 45 minutes before you can even identify the cause of action? AI-powered intake forms — built into your practice management system — can ask the right questions upfront, organise the facts chronologically, and hand you a pre-sorted brief before the client walks in. You still do the consultation. But you start at paragraph three instead of paragraph one.

Billing That Actually Gets Done

I've seen this a hundred times: advocates lose money not because they don't do the work, but because they're terrible at recording what they did. An AI agent connected to your billing system can take rough time entries — "worked on Kamau file, 3 hrs" — and turn them into properly descriptive narratives that clients understand and will pay. But notice: the AI needs access to the billing system. A standalone chatbot can't do this.

Contract Review and Risk Flagging

Not drafting the contract — reviewing it. Upload a 60-page lease or shareholders' agreement and ask AI to flag clauses that deviate from standard market terms, identify missing protections, or highlight ambiguities. You still read every line, as you must. But AI gives you a heat map of where to focus your deepest attention first.

Deadline and Limitation Tracking

AI connected to your case management system can analyse your open matters and flag which ones are approaching limitation periods, which have stalled without activity for 60 days, and which clients haven't been billed in three months. It's operational intelligence for the managing partner. But again — *it needs the system underneath.*

Compliance Monitoring

For firms advising corporates, AI can monitor regulatory changes — new gazette notices, CBK circulars, CMA guidelines, ODPC guidance — and flag what's relevant to each client's industry. Instead of manually scanning the Kenya Gazette every Friday, you get automated alerts matched to your client list.

Explaining Complex Documents to Clients

Kenya is a bilingual jurisdiction. Many clients, especially in conveyancing and family law, need explanations in Kiswahili or in plain language. AI is surprisingly good at taking a dense legal document and producing a client-friendly summary. It doesn't replace your advice — but it helps the client actually understand what you're telling them.

The Bottom Line

Here's what I'd want any managing partner to walk away understanding:

AI is not a product you buy. It's a capability you build into your operations. A chatbot subscription without a system underneath is like hiring a brilliant associate and then not giving them access to the file room, the library, or the billing software. They'll sit there and talk beautifully, but they can't actually do the work.

The firms that will win in the next five years aren't the ones with the fanciest AI tool. They're the ones that looked at every repetitive task in their practice — every form letter, every time entry, every compliance check, every status update — and asked: "Can a system handle this so my lawyers can think?"

And the most important thing you can do today, right now, this week? Give your people a regulated AI tool before their unregulated usage becomes your confidentiality crisis.

Tags: AI & Law

This article was originally published at:

<https://lenhac.com/blog/dear-partner-we-need-to-talk-about-ai>

Generated on June 15, 2026