AI agents inside the firm

What they can read, what they can file, and who signs

IN BRIEF

01 Law firms, accounting and audit firms, tax advisers and consultancies have started to put AI agents in front of their own people. One law firm has said that 3,400 lawyers and business services professionals will get an AI platform from January. Two of the largest accounting firms say they run agents by the hundred. Courts in the United States and England have already penalised filings that cited cases that do not exist.

02 The suppliers are moving in different directions. One platform for firms says in terms that on-premises deployment is not available. Two large accounting firms say their own agent platforms run inside the client's environment. A firm that sells advice needs to know which of these it is buying, and what the agent may do before a professional has looked at it.

03 This paper is about the work inside the firm: the reading, checking and drafting done at intake, in contracts, in transactions, in disputes, in audit, in tax, in proposals, in billing and in risk. It sets out what an agent can read, how it reaches a firm's systems, what it may change without a person, and what the courts and regulators already say. It ends with seven questions to put to any supplier, and with a place to start.

04 The position throughout is simple. An agent prepares. A named professional decides. Nothing is filed, signed or sent by the agent. The agent runs on infrastructure the firm controls.

IN THIS PAPER

01 Where law, accounting and consulting firms are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 6
06 A worked example: one mid-sized firm 8
07 Seven questions before an agent touches client files 9
08 Where to start 9

PART 01

Where law, accounting and consulting firms are today

3,400
lawyers and business services professionals one law firm said will get an AI platform

USD 5,000
penalty for filing opinions that did not exist, 22 June 2023

2313
cases in a public database of AI hallucination cases, read 22 November 2026

250
agents one large accounting firm is reported to have deployed across its own firm

Firms have started with the drafting. Addleshaw Goddard said on 9 November 2026 that "3,400 lawyers and business services professionals" will get access to an AI platform from January, and that it is keeping "human judgement, supervision and accountability at the centre" [1]. Trade press reports EY's agent platform as an initial 150 agents for 80,000 of its people [2], and PwC as having deployed 250 agents across its own firm [3].

The courts have noticed. In Mata v. Avianca, the United States District Court for the Southern District of New York found bad faith on the part of lawyers who filed non-existent opinions created by ChatGPT. On 22 June 2023 it imposed a penalty of USD 5,000 on the lawyers and their firm jointly, to be paid into the court [4]. In October 2025 a large consulting firm agreed to refund the last instalment of its fee for an Australian government report that held references to papers that do not exist and a fabricated quotation from a federal court judgment; the firm said the corrections "in no way impact or affect the substantive content" [5]. A public database kept by Damien Charlotin said "2313 cases identified so far", last updated on 20 November 2026, when we read it on 22 November [6].

On-premises platforms exist. EY's platform "will run across client clouds, on-premises, at the edge" [2]. PwC's agent platform "is deployed within the client's environment" [3]. Both are sold to clients, not handed to staff. The hosting pages we read for two products sold to firms name a public cloud: one says "Harvey hosts its cloud environment in Microsoft Azure", and the other says "On-premises deployment is not currently available" [7].

So where an agent runs is not the only question.

The harder ones are which tools it may call, as whom, with whose approval, and what record is kept.

What the published accounts leave out. None of the deployments we read describes a step where a named partner approves each document an agent prepares before it leaves the firm. The cases above involve people who filed text from a general chatbot without checking it. We found no public case caused by an agent that a firm built with approval gates, and we do not claim one.

One we can speak to. In production at a top-tier bank in Asia, air-gapped, inside the bank's own perimeter.

PART 02

The work nobody has packaged

One

template page we opened states a sign-off rule in words

None

was found for an audit sample, a tax return review or an independence check

We listed the jobs done inside a firm, from the conflict search to the partner's review of a bill, and looked for each in the template libraries of the agent and workflow platforms. The libraries cover the contract summary, the intake email and the time entry well. They thin out wherever a named professional signs something.

WHERE TEMPLATES EXIST	WHERE NONE WAS FOUND
Contract risk review and clause rewrites	**A sample drawn by the auditor's method, with how each item was picked**
Name matching against matter records for a conflict check	**A tax return set against the client's records and last year**
Time entry narrative rewrites	**A reply to a tax authority's letter, drafted for the adviser**
Engagement letters from an intake form	**An independence check prepared for the ethics partner**
Legal research from public case law	**An audit file read for missing sign-offs**
Invoices from timesheets	**Financial statements tied out to the trial balance**
	A pre-bill review that lists what the billing partner should look at

Two findings matter more than the list. First, review is rare and uneven. One template that writes legal analysis into a report and posts it to a chat channel describes no review step, though it calls its output "attorney-ready". One that rewrites time entry narratives writes them straight back to the record, and the page says nothing on whether anyone checks the flagged entries before billing. A contract review scenario from a large software supplier ends at "Summarizes findings for legal review" and says that it is "for demonstration purposes only". A conflict check template does add an email step where someone chooses to clear or escalate. Only one template page we opened states a rule in words: "Require legal-team review and sign-off before any contractual action" [8].

Second, the templates we opened are built from hosted services: a hosted model, a hosted chat service, a hosted spreadsheet or drive. That settles where client documents go before anyone has asked.

The gaps sit where the reading is slow, the records are controlled, and the decision already belongs to a named professional. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a firm that means three kinds of thing.

Documents.

Precedents, playbooks, the audit methodology, engagement terms, people policies, and the legislation, standards and case reports the firm holds. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the passage it relied on and says when the library is silent.

Systems of record.

Clients and matters, documents, time entries, bills, audit files, ledgers and signature records. Most are reached through an interface the firm already runs or its supplier publishes. The firm gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Data for analysis.

Time, billing, populations for sampling and trial balances, held in a data warehouse or reporting database, and read only.

All three pass one gate: every tool call is checked against the firm's policy, which can allow it, record it, hold it for a person, or block it.

Where the firm's interface checks the caller, the agent acts as the person using it and cannot claim another. Information barriers hold for the agent as they hold for the person.

SIX WAYS INTO A FIRM'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A contract, a client's letter, a set of accounts, a draft submission	Nothing is connected. The safest start
02	**A knowledge source**	Precedents, playbooks, methodology, engagement terms, legislation, standards	Permissions are kept per reader
03	**An imported description of an interface**	Matters, clients, time entries, bills, signature records	You choose the operations. Changes wait for approval
04	**A reporting database**	Time, billing, populations for sampling	Read only
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**Public registers and lists**	Company registers, sanctions lists, published legislation	Reads only, from the publisher

PART 03 · CONTINUED

148
interface descriptions taken from 15 publishers; the importer took 145

3,836
operations in the 145 that imported

1,788
of those operations read; 2,048 are sent as changes

714
of 1,258 operations in e-signature and agreement files change something

What is published. We took 148 interface descriptions from 15 publishers of the kinds of system a firm runs: e-signature and agreements, e-discovery, practice management, accounting ledgers and payroll, documents, time recording, contracts, proposals, consultancy project planning, a tax authority, sanctions lists and identity checks. The platform's importer took 145. It did not import 3, because three files use a notation it does not yet read. Together the 145 hold 3,836 operations [9].

What that exercise taught us. What firms' systems publish is built as much for changing as for reading. By method, 1,788 of the 3,836 operations read and 2,048 are sent as changes. The e-signature and agreement files alone carry 1,258 operations, and 714 of them change something: send for signature, void, create. Method alone cannot tell a read from a write, because some searches are sent as changes.

What is not published. We found no open description for the document management systems large firms run, for audit platforms, for tax compliance software, for legal research services, or for the conflicts and intake systems of large firms. Those come from the firm's own licence and tenant. Where nothing is published, the firm supplies its own description, or the agent reads from the data warehouse or from exports.

Client files and information barriers. The safe design keeps each matter's files where the firm keeps them, and applies the firm's own permissions and information barriers on every read.

So choose the operations at import, start with the ones that read, and add nothing that sends, signs or submits until the approval step has been seen working.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named professional decides.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the engagement letter, the markup, the chronology, the tax reply, the proposal, with each statement tied to its source.
03	**It lists** what is missing, what departs from the playbook, what does not agree.
04	**It stops** before anything that changes a record or reaches a client, a court, an authority or a colleague. A tool that changes something waits for approval.
05	**It decides nothing about a person.** It does not rank staff, reject a candidate or name anyone as the cause of an error.

Four lines are drawn tighter than the rest.

Courts, tax authorities and regulators.

No agent files, submits or serves anything. It may check every case, statute and quotation in a draft against the source in the firm's library and list each as found, not found, or not saying what the draft says. The professional who signs still reads the source.

Advice and opinions.

An agent may set out what the legislation, guidance and the firm's past advice say, with the passage. The advice to the client is the adviser's. No agent signs an opinion, an audit report or a tax return.

Conflicts and new clients.

An agent may list every possible match for a name, with the record it came from. It clears no conflict, accepts no client and approves no engagement.

People.

An agent may answer a colleague's question from the people policies, with the section cited. It may set each application against criteria the recruiter wrote and show the evidence. It does not rank, reject or assess a person.

PART 05

What the rules already ask

This part needs a lawyer's read before it is relied on. It reports what the texts say; it is not advice.

Courts and the duty to check

The duty is the lawyer's. In Ayinde v London Borough of Haringey and Al-Haroun v Qatar National Bank, the Divisional Court (Dame Victoria Sharp, President of the King's Bench Division, and Mr Justice Johnson) gave judgment on 6 June 2025. It said that "Freely available generative artificial intelligence tools, trained on a large language model such as ChatGPT are not capable of conducting reliable legal research", and that those who use AI to conduct legal research "have a professional duty therefore to check the accuracy of such research by reference to authoritative sources" [10].

Leaders are named. The court said that "practical and effective measures must now be taken by those within the legal profession with individual leadership responsibilities (such as heads of chambers and managing partners)", that "the profession can expect the court to inquire whether those leadership responsibilities have been fulfilled", and that guidance alone is "insufficient" [10].

Using AI is not the fault. In Mata, Judge Castel wrote that there is "nothing inherently improper about using a reliable artificial intelligence tool for assistance. But existing rules impose a gatekeeping role on attorneys to ensure the accuracy of their filings" [4].

PART 05 · CONTINUED

Agents, named

California speaks to agents directly. The State Bar of California's Practical Guidance, approved on 14 May 2026, replaces its 2023 version and addresses agentic AI. It says: "Attorneys must not deploy agentic systems in a manner that allows the system to make substantive legal determinations, communicate legal advice, prepare and file pleadings, or otherwise act in a representative capacity without meaningful attorney supervision and review." It is guidance for California lawyers, not a rule [11].

Client due diligence

The firm keeps the duty. Under the UK Money Laundering Regulations 2017, a firm may rely on a third party for customer due diligence, but "the relevant person remains liable for any failure to apply such measures" [12]. An agent that assembles identity and ownership documents leaves the duty with the firm. Whether an agent counts as a third party or an outsourcing provider under that regulation is a question for counsel.

The EU AI Act

The dates have moved. Regulation (EU) 2026/1744 set the high-risk duties to apply from "2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III" and from "2 August 2028" for those tied to Article 6(1) and Annex I [14].

A firm's drafting or review agent is generally not on the list. Annex III point 8 covers AI "intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law", or used in a similar way in alternative dispute resolution [13]. An agent that drafts, summarises or reviews for the firm's own lawyers is not that. It could be, if a firm's people used it for an arbitrator or a court. Which side a given use falls on is a question for counsel.

Recruiting is on the list. Annex III point 4(a) lists AI intended for recruitment or selection, "to analyse and filter job applications, and to evaluate candidates" [13]. A recruiting agent used that way is in scope, even where a recruiter decides. Counsel should read it before a firm turns one on.

People must be told, and the professional rules still apply. Article 50(1) asks that systems "intended to interact directly with natural persons" be built so that those persons "are informed that they are interacting with an AI system", unless that is obvious [13]. It would reach an intake assistant that talks to the public. The firm's own professional rules apply in any case.

Professional bodies

MARKET	WHAT WE READ	WHAT IT ASKS
United States	ABA Formal Opinion 512, 29 July 2024	Says the existing professional conduct rules apply to generative AI tools. As the Bar Examiner reports it, how much a lawyer must verify depends on the tool and the task, and managerial lawyers must set clear firm policies [15]
England and Wales	SRA warning notice "Misuse of AI", 17 August 2026	"The use of AI does not diminish or transfer your professional responsibilities." Reported by the Law Society of Ireland Gazette: 42 reports of potential misuse between July 2025 and July 2026 [16]
United Kingdom	FRC, guidance on AI in audit, 26 June 2025, and "Generative and Agentic AI Guidance", 30 March 2026	The first is "not prescriptive" and adds no new requirement. The second "discusses audit quality risks associated with GenAI and agentic AI tools, alongside possible mitigations" [17]

We did not open the primary text of the other bar and law society guidance, the audit standard setters' rules, the tax authorities' rules, or the rules in the Gulf, India or Singapore, and we make no claim about them. A firm in those markets should ask its own counsel. Privilege is a question for the firm's own lawyers; this paper does not say it is kept or lost.

PART 05 · CONTINUED

What the rules have in common

01	**The professional who signs stays responsible,** whatever tool prepared the work.
02	**Every authority a document cites is checked against its source** by a person.
03	**Client information is handled under the firm's own duties,** and with the client's consent where the rules ask for it.
04	**Leaders must set the policy** and be able to show it was followed.

PART 06

A worked example: one mid-sized firm

Take a made-up mid-sized firm with law, audit and tax teams in five offices: about 900 people. Nothing here is a firm's story, and every figure is illustrative; it is section 4 put to work.

01

At intake, the conflicts team gives the agent the names in a new instruction. The agent lists three possible matches, each with its record. The conflicts team clears two and refers one to the responsible partner, who declines the instruction.

02

Before a brief goes out, a lawyer uploads a draft with 31 citations. The agent finds 29 in the firm's library and shows the passage for each. One is not found, and one passage does not say what the draft says. The lawyer reads both sources and corrects the draft before signing.

03

In audit, the agent reads a file against the methodology and lists 11 working papers with no sign-off and 4 open review notes. The engagement manager chases them. The partner concludes.

04

In tax, a letter arrives from a tax authority. The agent sets out what it asks and the date, and drafts a reply from the client file. The adviser rewrites two paragraphs and sends it.

05

In billing, the agent sets a draft bill against the engagement letter and the time recorded, and lists six entries with vague narratives. The billing partner decides every write-off.

06

Before an audit engagement is accepted, the agent sets the proposal against the list of audit clients and the services already supplied, and lists what the independence rules say to consider. The ethics partner decides.

Every one of these reads documents or records the firm already has, files nothing, changes no record by itself, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches client files

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? Can it run with no connection to the internet?

02 **What can it read?**
Does it act as the person using it, with their permissions and our information barriers, or as one all-seeing account?

03 **What can it change?**
Which tools change a matter record, a time entry, a bill or a letter, and does each wait for a person?

04 **Can it reach a court, a regulator or a client?**
Show us that no tool files, submits or sends, and who presses the button.

05 **Who decides, and who reads the source?**
Show us the step where a named professional approves, and where each citation appears with its passage.

06 **What is recorded?**
Who asked, which tool ran, what was decided, and for how long is that kept?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a professional, and no client, court or authority waits on the agent.

Citation checks
Every case and quotation in a draft is set against the source before the lawyer signs. Nothing leaves the firm.

Conflict check preparation
The agent lists every possible match with its record. The conflicts team decides.

Audit file review
The agent lists missing sign-offs and open review notes before the partner's review. The partner concludes.

Each of the three reads documents the firm already holds. None reaches a client, a court or an authority, and in each the person who signs today is the person who signs tomorrow. That is why they can start before any system of record is connected.

One team, one office, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/PROFESSIONAL-SERVICES

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 Harvey blog, "Addleshaw Goddard and Harvey launch firmwide collaboration to shape AI-enabled legal delivery", 9 November 2026. harvey.ai

02 Accounting Today, article on the EY and Deloitte agentic AI partnership with Nvidia, 19 March 2025. accountingtoday.com

03 Accounting Today, "PwC announces agent OS for coordinating AI agents", 27 March 2025. accountingtoday.com

04 United States District Court, Southern District of New York, Mata v. Avianca, Inc., 22-cv-1461 (PKC), Opinion and Order on Sanctions, 22 June 2023 (Document 54). courtlistener.com

05 Fortune, article on a consulting firm's refund over an Australian government report, 7 October 2025. fortune.com

06 Damien Charlotin, AI Hallucination Cases database, last updated 20 November 2026, read 22 November 2026. damiencharlotin.com

07 Harvey security page, harvey.ai/security; Intapp, Celeste governed AI page, intapp.com.

08 Template libraries of agent and workflow platforms, searched 22 November 2026: n8n templates 18250, 14434, 17708 and 19289 (n8n.io); Box support article, "Contract review agent" (support.box.com); Microsoft adoption scenario library, "Automated contract review agent" (adoption.microsoft.com).

09 148 interface descriptions fetched from their publishers on 22 November 2026 and run through the platform's importer (Katonic research).

10 High Court of Justice, King's Bench Division, Divisional Court, R (Ayinde) v London Borough of Haringey and Al-Haroun v Qatar National Bank [2025] EWHC 1383 (Admin), 6 June 2025, paragraphs 6, 7, 9 and 82. caselaw.nationalarchives.gov.uk

11 State Bar of California, Practical Guidance for the Use of Generative AI in the Practice of Law, approved 14 May 2026. calbar.ca.gov

12 The Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017, regulation 39(1). legislation.gov.uk

13 Regulation (EU) 2024/1689 (AI Act), Annex III points 4(a) and 8(a); Article 50(1).

14 Regulation (EU) 2026/1744, 8 July 2026, amended Article 113.

15 American Bar Association, Formal Opinion 512, "Generative Artificial Intelligence Tools", 29 July 2024, as reported in The Bar Examiner, Fall 2024. thebarexaminer.ncbex.org

16 Solicitors Regulation Authority, warning notice "Misuse of AI", 17 August 2026. sra.org.uk; Law Society of Ireland Gazette, "SRA issues warning notice on use of AI", 19 August 2026. lawsociety.ie

17 Financial Reporting Council, news release on its AI in audit guidance, 26 June 2025; AI in audit guidance page listing "Generative and Agentic AI Guidance", 30 March 2026. frc.org.uk