MAKES ELEVEN CONSULTING · GUIDEBOOK

The managing partner's guide to AI in a law firm

For the partner whose name is on the answer when the committee, the client, or the carrier asks how AI is controlled here.

Most guides of this kind are written for people deciding whether AI is interesting. This one is written for the person who already knows it is here, has been handed responsibility for it, and now has to say something defensible to partners who disagree with each other. It maps the duties an ethics opinion actually names to things a firm can be asked to produce, sets out the three ways firms get into trouble, and ends with a checklist your committee can run against any vendor, this one included.

Written for
Managing partners and AI-committee chairs at US firms of roughly 10 to 150 attorneys.

Mapped to
ABA Formal Opinion 512, duty by duty, with the operational demand under each one.

Sourced
Every external claim carries a citation at the end. Anything we could not source was cut.

Makes Eleven Consulting is a boutique advisory firm that uses AI. It does not practice law, and nothing in this document is legal advice or an ethics opinion. Your jurisdiction's rules govern; your ethics counsel reads them.

START HERE

What this is, and who it is for.

Someone at your firm has been made responsible for AI. In most firms this size it was not a job anyone applied for. It arrived because a partner started using a chatbot on a brief, or a client sent an outside-counsel guideline with an AI clause in it, or an associate asked a question at a partners' meeting and nobody had an answer. If that person is you, this is written for you.

The hard part of the job is not choosing software. It is that you are answerable to four audiences who want different things. Partners want permission or prohibition and will accept either, as long as it is clear. Clients want to know their material is not going anywhere. The carrier wants a policy with a name attached to it. And the associates want to know whether the firm is investing in them or replacing them. A tool decision answers none of those. A written standard answers all four, which is why this document spends its pages there.

What you get out of this

Five questions you are going to be asked, with answers that hold up when pushed. The duties ABA Formal Opinion 512 names, each translated from a principle into a thing a firm can be asked to produce. The three failure modes that actually generate discipline, and what stops each one. A boundary between work that can be automated today and work that must not be. A checklist your committee can run against any vendor. And the shape of a first engagement, if you decide you want one.

What this is not

It is not legal advice, and it is not an ethics opinion. We are not lawyers and we do not practice law. Where a question turns on your state's rules, the honest answer is that your ethics counsel reads them and we do not. It is also not a product survey. There is no tool comparison in here, deliberately, because the tool is the last decision a firm should make and the one most firms make first.

START HERE · CONTINUED

Why we wrote it, and how to use it.

Why we wrote it

Legal is a named practice at this firm, not a detour into an adjacent market. We publish the standard we build to, we publish the list of actions our systems refuse, and both are listed at the end as companion documents. This guidebook is the argument that sits above them. If you read it and conclude your firm needs a policy and one careful hire rather than anything we sell, that is a legitimate outcome and we would rather you reach it here than after an engagement.

How to use it with your committee

Take the six duties one at a time and answer only this: what could the firm produce this month if someone asked. Where the answer is nothing, you have found an agenda item. Then read the fifteen committee questions aloud and mark the ones nobody in the room can answer about software the firm already uses.

A note on sourcing

This document names one court decision, one ethics opinion, one local rule, and one public database. All four are cited in full at the end with the date each was checked. Several claims that would have made it more dramatic were cut for want of a source you could open.

THE ROOM

Five questions you will be asked, and answers that hold.

Not the questions a vendor prepares you for. The ones that arrive without warning, from people entitled to an answer.

01 "Is this going to get one of us disciplined?"

Asked by a partner, usually a senior one, usually in a meeting about something else.

The honest answer is that the risk is real, well documented, and almost entirely procedural. Lawyers have been sanctioned for filing briefs citing cases that do not exist. What produced those sanctions was not the model. It was that no step between the draft and the filing required a citation to resolve to a real document. Firms that have added that step have not generated these cases. So the answer to give the partner is narrow and true: the exposure comes from unverified output reaching a court, and the firm controls that gate directly, today, without buying anything.

02 "Is our client material going into a model?"

Asked by a client, an insurer, or a general counsel with an outside-counsel guideline in hand.

This one has a factual answer, and most firms cannot give it. The question is not whether AI is used. It is which tools retain what is typed into them, under what contract, and who at the firm is permitted to put matter material into which. Opinion 512 is direct about the mechanism: a self-learning tool can surface one client's information in another user's session. The answer that satisfies a sophisticated client is a short list of approved tools, the retention terms for each, and the rule about what may be entered. If you cannot produce that list this week, that gap is the finding, and it is a week of work rather than a project.

03 "Does the firm have a written AI policy, and who enforces it?"

Asked by the malpractice carrier, or by the client's procurement team, in writing, with a box to check.

We cannot tell you what your carrier asks, and we would be making it up if we tried. We can tell you the shape of an answer that survives follow-up. It has a named seat rather than a committee, because a committee cannot stop a workflow at four on a Friday. It has a list of approved uses and a list of prohibited ones, and the prohibited list is the more useful document. It has a record of who was trained and when. And it is dated, because an undated policy tells a reader it was written for an application form.

What the first three share

Not one of them is answered by naming a piece of software. Each is answered by a short document: a list of approved tools with their retention terms, a dated policy with a named seat, and a training record. None of the three takes longer than two weeks to produce, and a firm that has them can answer a client, a carrier, and a partner from the same folder.

THE ROOM · CONTINUED

The two nobody prepares for.

04 "What are we actually getting for this?"

Asked by the partner who votes on the budget, usually about eight months after the first tool was bought.

This is the question that ends most AI programs, and it is fair. The pattern is familiar enough to be predictable: a tool gets adopted, everyone reports that it helps, and no line on the firm's financials moves. That happens because a fast tool was attached to a process whose speed is set somewhere else. If a memo is drafted in twenty minutes instead of four hours but still waits nine days for a partner to open the folder, cycle time did not change and neither did anything you can bill or staff against.

The answer is to pick one number before you start, record it on day one, and report against it. Matters closed per month. Days from intake to demand letter. Hours between a document landing and it being indexed. Pick the number the managing partner already watches, not a new one invented for the pilot. If it does not move, the honest report says so and the pilot stops. That is a result, and it costs far less than a program nobody is willing to end.

05 "Am I being replaced?"

Asked by an associate or a paralegal, rarely out loud, and answered by what the firm does rather than what it says.

The distinction that actually holds in a law firm is between assembly and judgment. Assembly is finding the document, building the chronology, reconciling the ledger, noticing the deadline, pulling the same four exhibits a fourth time. Judgment is deciding what any of it means and putting a name to that decision. Systems are good at the first and are not permitted the second, and in a law firm that is not a philosophical position. It is a licensing one.

The practical answer is that the work being automated is the work that trains nobody. A second-year who spends a week Bates-stamping learns less than one who spends the same week reading a record that has already been assembled. Say that out loud, and then check whether the firm's pilot actually targets that work. If the first project touches judgment rather than assembly, the associates were right to worry and the committee picked wrong.

WHO IS ASKING, AND WHAT THEY WILL ACCEPT

Partners	Clients and carriers	The budget vote	Associates
A clear rule about what is allowed, and a name attached to it.	The approved-tool list and its retention terms, in writing.	One number, recorded on day one, reported against.	A pilot that visibly targets assembly and not judgment.

THE STANDARD

ABA Formal Opinion 512, duty by duty.

Issued 29 July 2024 by the ABA Standing Committee on Ethics and Professional Responsibility.

Opinion 512 binds nobody. The Model Rules are a model, and your jurisdiction's rules and its own opinions are what govern your firm. What 512 is good for is structure. It applies rules your partners already know to a technology they do not. Its opening sentence lists the obligations a lawyer using these tools has to consider, and they gather under six headings, which makes a serviceable agenda for a committee that does not know where to start. The duty to advance only meritorious claims is grouped with candor below, since in practice the two fail together.

The gap most committees fall into is between reading the duty and discharging it. A duty is a principle. What a regulator, a client, or an opposing party asks for is an artifact. So each duty below is paired with the thing your firm should be able to hand over without preparing it first.

One thing 512 does not do

It does not tell you which tools are acceptable, and it declines to set a fixed amount of verification. Both are left to judgment about the specific task, which means a firm cannot discharge these duties by buying a product. The product is a component. The standard is the firm's.

Competence

Model Rule 1.1

What the opinion says. A lawyer need not become an expert in the technology, but must have a reasonable understanding of what a specific tool can and cannot do. The opinion is pointed about uncritical reliance on generated content, and it treats the amount of verification required as varying with the task rather than being uniform.

What you must be able to show. A named list of approved tools, and for each one the person who evaluated it, what they tested, and what the tool is not approved for. "Some of us use a chatbot" does not discharge this duty, because nobody at the firm can say what was relied on or how far.

Confidentiality

Model Rule 1.6

What the opinion says. Self-learning tools raise the risk that information from one representation is disclosed in another user's session. Where client information is entered into such a tool, the opinion calls for the client's informed consent, and says that general boilerplate added to an engagement letter is not sufficient: the client needs the lawyer's judgment about why the tool is used, what the specific risks are, and what will be disclosed.

What you must be able to show. Which tools retain input and on what contractual terms, who may enter matter material into which, and where consent was needed, a record of it that names the tool and the risk rather than gesturing at "technology." This is the duty that most rewards saying no to consumer tools early.

THE STANDARD · CONTINUED

Supervision, communication, and candor.

Supervision

Model Rules 5.1 and 5.3

What the opinion says. Managerial and supervisory lawyers must establish clear policies on the use of these tools and make reasonable efforts to see that the firm complies. Training is expected to cover both the tools and the ethical obligations attached to them. The duty reaches outside the firm: an outside GAI provider is vetted the way earlier opinions vetted any outside vendor, down to its security policies, its conflicts screening, notice of a breach, what it retains when the engagement ends, and the forum available if the agreement is violated.

What you must be able to show. A policy that predates the tool, a training record with names and dates, and one seat that can stop a workflow without convening anybody. The test of a supervision answer is whether you can name the person, not the committee.

Communication

Model Rule 1.4

What the opinion says. Telling a client about AI use depends on the circumstances. Disclosure is called for when the client asks, when client information is entered into the tool, and when the use bears on the basis of the fee. Consultation is also necessary when the tool's output will influence a significant decision in the representation; the opinion's examples are evaluating litigation outcomes and jury selection. The engagement agreement is an appropriate place to address it.

What you must be able to show. A written decision about what your engagement letter says and what matter opening asks, plus the ability to answer "was AI used on my matter" from a record rather than from memory. The second half is the one firms discover late.

Candor toward the tribunal

Model Rules 3.3 and 3.1

What the opinion says. Output must be verified before it is relied on: the opinion names citations to nonexistent opinions among the issues that have already arisen, and holds that even an unintentional misstatement to a court can involve a misrepresentation. Courts have sanctioned lawyers for fabricated citations, but that is case law rather than the opinion; the first failure mode covers it. The duty to advance only meritorious claims sits alongside it.

What you must be able to show. Verification as a step somebody performs and initials, not a sentence in a policy. And a per-filing answer to the local rules, which now reach this directly: in the Northern District of Texas, a brief prepared using generative AI must say so on its first page, and a brief filed without that disclosure certifies that no part of it was so prepared. Silence is an affirmative representation. Your firm needs to know, filing by filing, which of those two statements is true.

THE STANDARD · CONTINUED

Fees.

Fees

Model Rule 1.5

What the opinion says. On hourly work, a lawyer bills the time actually spent; efficiency the tool produces does not convert into billable hours. Flat fees answer to reasonableness instead: charging the same flat fee with the tool as without it may be unreasonable, because "[a] fee charged for which little or no work was performed is an unreasonable fee." The cost of a tool may be overhead or a chargeable expense, depending on the tool.

What you must be able to show. A recorded decision, made before any of this reaches a bill, about how AI-assisted time is captured and whether tool cost passes through. Firms that leave this to individual judgment discover the answer during a fee dispute, which is the worst available time to form one.

THE THREE FAILURE MODES

What actually produces discipline, and what stops it.

Three ways firms get hurt. None of them is the technology behaving strangely.

One · The fabricated citation

In June 2023 a federal judge in the Southern District of New York sanctioned two lawyers and their firm over an affirmation citing six judicial decisions that did not exist. The citations had case names, reporter volumes, and page numbers. When the fabrications were questioned, the lawyer went back to the same tool and asked whether the cases were real, and it said they were. That detail is the whole lesson. The system was not asked to produce a source; it was asked for reassurance, and it obliged.

It is no longer an isolated event. A public database maintained by a researcher at HEC Paris tracks decisions where a court found or implied that a party relied on hallucinated material, plus, by its own note, a small number where the AI use was alleged but not confirmed. It recorded 1,981 such decisions worldwide when we checked it on 29 August 2026, 1,363 of them in the United States.

The control. Sourced answers or a refusal. A citation is emitted only when it resolves to a record that can be opened, and below the confidence threshold the output is a question to a human rather than a sentence. That is an engineering property rather than an instruction: the refusal to assert without a source, and the refusal to guess below threshold, are two of the enumerated classes in our prohibited-actions register. A firm can adopt the same rule without buying anything, by making a citation check a step someone signs rather than a habit someone has.

Two · The confidentiality leak

This one rarely looks like a breach. It looks like ordinary work. A lawyer pastes a section of a deposition into a consumer tool to get a summary, and the tool retains it under terms nobody at the firm has read. Opinion 512 names the mechanism plainly: material from one representation can surface later in a different user's session, and be relied on, filed, or repeated to someone else.

There is a second cause that gets less attention, and it catches firms that did everything else right. Many AI systems log prompts and outputs for debugging. Those logs are a second copy of privileged material, held wherever the vendor keeps them, and no retention schedule at your firm covers them. It surfaces during a discovery dispute or an exam, which is late.

The control. Per-client isolation of credentials, storage, and memory, and a no-training clause in your own contract rather than a marketing line. Where the pattern fits, the work runs on the firm's own subscription and the firm's own storage, so the party holding privileged material is the firm, as it was before any of this. And run records name their sources as identifiers rather than reproducing contents, so the log does not become the second file.

THE THREE FAILURE MODES · CONTINUED

The one that comes from a calendar, not a model.

Three · The unreviewed filing

Nobody decides to skip review. What happens is that review was never anybody's named job, so it lived in the space between the drafter and the filer. On an ordinary Tuesday that space is filled by a careful person. On the Friday before a deadline, with the responsible partner in a mediation, it is filled by nothing at all. The document goes out because the calendar says it must, and the control that existed on paper was a policy rather than a gate.

The distinction matters and it is easy to test. A policy is a sentence that describes what people should do. A gate is a step the work cannot get past. If your review requirement can be satisfied by someone intending to have read something, you have a policy. If the filing cannot leave without a second name attached to it, you have a gate.

The control. Nothing unreviewed reaches a court, a regulator, or a counterparty, and no system we build signs or files anything. Both are enumerated refusals rather than assurances, which means the check is that the system cannot perform the action, not that it has been asked not to. In your own firm the equivalent is cheap: name the seat, and make the last step before a filing a person rather than a process.

WHAT THE THREE HAVE IN COMMON

None of them is a technology failure.

Read the three together and the pattern is hard to miss. A model produced text, which is what it does. In each case the firm had no step that required the text to resolve to something real, no rule about where material could be typed, and no person whose name was attached to the last look. Those are process defects, and they were process defects before AI existed. What changed is the volume of plausible output arriving at the gate. A gate that was adequate at four documents a week is not adequate at forty, and the firms getting hurt are usually the ones whose controls were informal and worked fine for twenty years.

The fabricated citation
Fails because nothing required a source. Fix the verification step.

The confidentiality leak
Fails because nothing said where material may go. Fix the approved-tool list.

The unreviewed filing
Fails because nothing named the reviewer. Fix the seat.

THE BOUNDARY

What to pilot first, and what must never run unattended.

One test does most of the work: can the output be checked against a source, and does the step end in a signature?

Assembly	START HERE
Drafting	ONLY BEHIND A NAMED REVIEWER
Judgment · signature · money	NEVER

Start here · assembly

Work that is high in volume, verifiable against a record, and ends with nobody signing anything. Document intake and indexing. Chronologies built from a produced record, with every entry citing its page. Deadline and docket monitoring. First-pass responsiveness and privilege sorting, with a human deciding every call the system flags. Conflicts intake. Trust reconciliation and pre-bill hygiene, where every figure is recomputed independently and exceptions are listed rather than quietly corrected. This band is where the hours are, it is where the errors are checkable, and it is the least glamorous work in the building, which is why it gets skipped for something that demonstrates better.

Only behind a named reviewer · drafting

First drafts of routine correspondence, discovery responses, and research memoranda. There is nothing wrong with the draft. The control is the release, and it has to be a person with the authority and the time to reject the thing. If the reviewer is a partner with a trial next week, the reviewer is decorative and the band has collapsed into the one above it.

Never · judgment, signature, money

Advice to a client. Case strategy. The decision to file, and the signature on the filing. Settlement authority. Any movement of money, client trust funds most of all. Anything that reaches a court, a regulator, or an opposing party without a human having read it. These are not on the list because the technology is weak. They are on it because each one is an act a licensed person performs and answers for, and delegating it does not transfer the answering.

The line your committee can adopt

Assembly is not judgment. Our systems assemble the file, the chronology, and the reconciliation. The attorney owns every judgment, every signature, and every dollar. No system we build signs, files, or pays, and that is written into the engagement rather than implied by it.

FOR THE COMMITTEE

Fifteen questions to ask any vendor, this one included.

Written to be read aloud in a meeting. The failure signal is a vendor answering a different question than the one asked.

01 **Where does our data sit, and who holds the keys?**
A good answer names infrastructure. A bad one describes how seriously they take security.

02 **Does anything we enter train a model? Show us the clause.**
The clause, not the marketing page. If it is only on the website, it is not a term.

03 **What does the system do when it does not know?**
"It is very accurate" is a failed answer. You are asking for the refusal behavior, and there either is one or there is not.

04 **Who checks the work, and did they build it?**
If the builder is the examiner, there is no examination. Ask who the second party is and what they were given to test against.

05 **What can this system do that we have not approved?**
Ask for the list of actions it refuses. A vendor who has never written one is telling you something.

06 **What is logged, and does the log hold client content?**
The log is the second copy of the file. Ask where it lives and how long it is kept.

07 **What happens when we end this? What do we migrate?**
The best answer is nothing, because everything ran on your own subscription and your own storage.

08 **Who is accountable, by name, on day one?**
A named person, in writing, at the start. Not an account team.

09 **What is the ceiling, in writing, for the term?**
Scope and a ceiling agreed before work begins. Anyone who cannot commit to a ceiling is planning to exceed one.

10 **What is the baseline, and what will we see at day 30?**
A number recorded before the work starts, and a before-and-after against it. Otherwise the review is a conversation about vibes.

11 **When is the answer "ask your ethics counsel"?**
A vendor who never says this is either unaware of the line or willing to cross it.

12 **What would you refuse to build for us?**
The most useful question on this page. An answer with nothing in it means the refusals will be discovered during the engagement.

13 **How would you know if you were adverse to one of our clients?**
Opinion 512 carries the old vendor-vetting duties over to GAI providers, conflicts screening among them. A provider who has never been asked this has never screened.

14 **When are we notified of a breach, or of legal process seeking our data? Show us the clause.**

Notice of a breach or of service of process is a contract term or it does not exist. Here too, the marketing page is not a term.

15 **What do you retain after we leave, what rights do you claim in it, and in what forum could we seek relief?**

Retention after discontinuation, proprietary claims on what was entered, and an accessible legal forum are points the opinion tells a firm to pin down in the vendor terms.

Ask us all fifteen. Our answers to most of them are already published: the written standard is on the assurance page at makeseleven.com, and the full enumeration of what our systems refuse is one of the companion documents on the next page.

IF YOU WANT HELP

Where a firm typically starts.

One workflow, one accountable seat, one month running both ways and timing them.

The first thing is a diagnosis, and it is deliberately unglamorous. We map how a piece of work actually moves through your firm, which is almost never how the process document says it moves, and we find where the time goes. The output is a written picture of the current state and the outcome that matters, and you keep it whether or not anything follows. Firms that stop there have still gained the thing they were missing, which is an accurate description of their own process.

What follows, when something does, runs on the same five moves every time: Audit, Gauge, Engineer, Navigate, Track. The companion short read sets them out in full and names the artifact each stage leaves behind. The published sprint is roughly eight weeks, in three phases, and the shape exists to force a visible result while the people who will live with the system are still paying attention.

Weeks 1–2

Audit & Gauge

Map the workflow as it runs, find where the time actually goes, and choose the one number that will decide whether this worked.

Weeks 3–5

Engineer

Build and deploy the first workflow. Guardrails ship with it rather than after it.

Weeks 6–8

Engineer · Navigate · Track

Scale what worked, wire the oversight, and prove the number against the day-one baseline. Then choose again.

Two things worth saying before you call

The proposal contains only what the diagnosis supports. If the honest finding is that your firm needs a written policy, an approved-tool list, and one careful hire, that is the report you get, and it is a shorter engagement than the one you were expecting. Scope and duration are set per engagement, because a twelve-lawyer practice and a hundred-and-forty-lawyer firm are never the same shape.

And the boundaries are contractual rather than cultural. Every financial deliverable is inspected by a licensed CPA in our delivery network before it reaches you, disclosed in the engagement letter; this is advisory work, and no financial-statement audit, review or other attest report is issued. We do not practice law and nothing we build gives legal advice. No system we build signs, files, or pays. Those are terms you can hold us to on paper.

What the first conversation is

Thirty minutes, free, and it does not need a committee. Name the task that eats the most hours in your practice and you will leave the call knowing whether it is work a system should touch at all. Sometimes the answer is that it is not, and hearing that costs you half an hour.

NEXT

What to read, what to try, and the one conversation worth booking.

The two companion documents

The A.G.E.N.T. method, a short read. Three pages on the five moves every engagement runs, the four questions that come before them, the engine underneath, and the artifact you keep at the end of each stage.

The prohibited-actions register. The enumeration of every class of action our systems are built to refuse, each with the basis for the refusal and how it is enforced. It is the document to read if you want to check whether a vendor's boundaries are real, because it shows what one looks like when written down properly.

Two things that are free, permanently

The Workflow Mapper. Ten minutes, no account, no email. Map one process and see where it actually stalls. Nothing you build leaves your own browser unless you deliberately send it to us. At makeseleven.com/map.

The Workflow Atlas. One board in every vertical is open to everyone with no trial clock. The law vertical holds nine boards, eight practice areas and a case-evaluation board, each mapped step by step with the autonomy ceiling and the seat that can stop it; the case-evaluation board is the open one. At makeseleven.com/workflow-atlas.

THE ONE CALL TO MAKE

The AI Committee Briefing.

Fifty minutes for the whole committee. The fears named accurately, the boundary between assembly and judgment, the tier ladder, the named seat, and what your firm can do on Monday. It is enablement rather than a pitch, and it books directly with no diagnosis required. If it ends with your committee writing its own policy and never speaking to us again, the briefing did its job.

Book it	Ask a question first
makeseleven.com/consultations	**leonfields@makeseleven.com**

Makes Eleven Consulting is a boutique advisory firm that uses AI, working nationwide, founded and run by Leon Fields II. Legal and policy is a named practice with its own workflow boards, its own four-quarter curriculum, and a productized desk for the accounting a trust exam checks. Every maintained engagement carries 24/7 support through the Help Desk at makeseleven.com/help-desk.

SOURCES

Everything cited, with the date we checked it.

Four external sources carry every factual claim in this document. All four were opened and read on 29 August 2026.

ABA Formal Opinion 512

American Bar Association Standing Committee on Ethics and Professional Responsibility, *Generative Artificial Intelligence Tools*, Formal Opinion 512, issued 29 July 2024.

americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf

Mata v. Avianca, Inc.

678 F. Supp. 3d 443 (S.D.N.Y. 2023), No. 22-cv-1461 (PKC), signed 22 June 2023. The opinion sanctioning two lawyers and their firm over six fabricated citations.

law.berkeley.edu/wp-content/uploads/archive/2025/12/Mata-v-Avianca-Inc.pdf

AI Hallucination Cases

Database maintained by Damien Charlotin, HEC Paris. Tracks decisions where a court or tribunal found or implied that a party relied on hallucinated material; by its own note it also includes some decisions where AI use was alleged but not confirmed. 1,981 decisions worldwide, 1,363 in the United States; database last updated 28 August 2026.

damiencharlotin.com/hallucinations/

N.D. Tex. Local Civil Rule 7.2(f)

United States District Court for the Northern District of Texas, Local Civil Rules, effective 2 September 2025. Rule 7.2(f) requires first-page disclosure where a brief was prepared using generative AI, and provides that a brief filed without the disclosure certifies that no part of it was so prepared.

txnd.uscourts.gov/sites/default/files/documents/CIVRULES.pdf

On the method

The A.G.E.N.T. method, its agent taxonomy, and its sprint structure are adapted with attribution from Hofmann, D. and Kruhse-Lehtonen, U., "The Agent-Centric Enterprise: Why 2–10x Productivity Gains Demand Radical Workflow Redesign," *Harvard Data Science Review*, issue 8.1, winter 2026, at hdsr.mitpress.mit.edu. What this firm adds is the assurance architecture around it. Public-sector and framework alignment work follows the NIST AI Risk Management Framework, at nist.gov/itl/ai-risk-management-framework.

What we left out

Several things a document like this usually contains are missing on purpose. There are no adoption statistics, because the ones in circulation are survey figures we cannot independently check. There is no claim about what malpractice carriers now require, because we do not know and neither does anyone quoting a number at you. There are no client names and no case outcomes. And there are no prices, because scope and a written ceiling are set per firm, in a proposal, after somebody has looked at the actual work.