CLEMENZA LAW GROUP

NEW YORK, NY · (212) 365-4875

THE PRIVATE BRIEF · HANDBOOK

The AI at Work Handbook

Somebody in your organization is using an AI tool right now. Possibly several people. They are summarizing documents, drafting emails, tidying reports, answering customer questions, and occasionally making decisions with it. Some are using a tool the company pays for. Some are using a free account on their phone. Almost nobody has been told what they may and may not paste in, and almost no board has discussed any of it.

That is not a scandal. It is the normal state of most companies and nonprofits in this decade, and it is entirely fixable in a few afternoons. This handbook is not an argument against these tools; we build with them ourselves. It is about knowing four things: what goes into them, what comes out, what you have promised other people, and who is responsible. Once you know those, the tools stop being a quiet risk and become what they should be, which is useful.

Read it once now. Then keep it for the day a customer asks how you use AI, a vendor changes its terms, or a board member asks the question nobody has an answer to.

clemenzalaw.com/reference/ai-at-work

Chapters

CHAPTER 1

Before you begin: what actually goes wrong

Start with the failure that actually happens, because it is not the one people imagine. It is not a rogue algorithm or a dramatic leak. It is a capable employee at six in the evening, under deadline, pasting a customer's contract into a free chatbot to get a summary, because it was faster and nobody had ever told them not to. Nothing bad happens that day. The problem is what the tool's terms allowed it to do with that contract, and what the customer's contract said about where their information could go, and the fact that nobody in the company knows either.

Every AI problem an organization is likely to meet falls into one of three exposures. What goes in: confidential information, customer data, personal information protected by privacy law, material you licensed but do not own, entered into a system whose terms you have not read. What comes out: text, images and decisions that may be wrong, may be biased, may reproduce someone else's work, and may not be protected by copyright at all when you want to enforce it. And what you promised: warranties in customer contracts written for a world where people made the deliverables, marketing claims about what your AI does, notices you owe job candidates, and statements to investors.

Each exposure has a small, specific fix, and none of the fixes require slowing anyone down. They require knowing, in writing, what tools are in use, what may go into them, what a person checks before the output leaves the building, and who owns the question. The chapters that follow are those fixes in order. How long the work takes depends on how many tools are in use, what data they touch, and what your customer contracts already promised.

One idea runs through the whole handbook and is worth stating once. A named person has to review what goes out and answer for it. "The tool wrote it" has never been a defense anywhere it has been tried, not in a court, not with a regulator, not with a customer. That single rule, understood by everyone, prevents most of what follows.

THE CHECKLIST

- [ ] **Name the three exposures out loud: what goes in, what comes out, what you promised**

  They are the whole map. When a new tool, a new use or a new question arrives, ask which of the three it touches and the answer tells you what to check.

- [ ] **Adopt the one rule now, before the policy exists**

  A named person reviews it and answers for it. Every factual claim, number, name and citation in anything that leaves the building has been checked by a human who is accountable for it.

  > This is the rule that would have prevented the lawyers sanctioned for filing invented cases, the marketing team that published a competitor's paragraph, and the analyst whose confident wrong number reached the board.

- [ ] **Decide that this is a governance matter, not an IT matter**

  The risks are legal, contractual and reputational. The person who owns them should be able to read a contract and say no, which is usually an operations lead or a general counsel, not the most technical person in the room.

CHAPTER 2

Find out what is already in use

You cannot write rules for tools you do not know about, and every organization we have looked at has more in use than anyone thought. The inventory is the first job and it takes less time than it sounds.

There are three kinds of AI in a typical organization. The tools people chose: the chatbots and assistants employees use directly, on company accounts or personal ones. The features vendors turned on: the AI summaries inside your email, your customer-service platform, your document system, your recruiting software, often switched on by the vendor in an update nobody read. And the systems that make or shape decisions: pricing, screening, scoring, routing, anything that sorts people or money automatically. The third kind is the smallest list and the one with the most law attached.

The fastest way to build the inventory is to ask. Ten minutes at an all-hands, with one honest promise: nobody is in trouble for the tool they have already been using. That promise is what turns shadow use into disclosed use. You will learn more in those ten minutes than an audit would tell you, and the list will be longer than the approved one. That gap is where the policy starts.

A board does not need a technical document. It needs a one-page table with three columns: what, where, and who owns it. Most boards are surprised by its length and by how much of it was never approved by anyone.

THE CHECKLIST

- [ ] **Ask everyone what they use, with amnesty**

  One meeting, one question, one promise. Write down every tool, the account it runs on (company or personal, free or paid), and what it is used for.

  > Without the amnesty you get the approved list back. With it you get the real one. The difference is the whole point.

- [ ] **Check what your existing vendors have switched on**

  Email, documents, customer service, recruiting, accounting, analytics. Open the settings. Many have added AI features by default, sometimes with data terms different from the rest of the product.

- [ ] **List separately anything that decides or scores**

  Hiring screens, credit or eligibility decisions, pricing, fraud flags, content moderation. These carry specific legal duties (see the chapter on decisions about people) and need a human in the loop.

- [ ] **Put it in a three-column table: what, where, who owns it**

  This is the document the board sees, the insurer asks for, and the policy is written against. Keep it current; review it twice a year.

- [ ] **Note which accounts are personal**

  A consumer account and the business tier of the same product are different products with different data terms. Personal accounts used for work are the first thing the policy will address.

CHAPTER 3

Write the two-page rule

A company of twenty, or a nonprofit of fifty, does not need an AI governance framework. It needs a short document that answers four questions clearly enough that an employee under pressure can follow it without asking anyone. Which tools are approved. What data may go into them. Which uses are permitted and which are not. And what a person must check before the output leaves the building. Add one named owner and one review date, and you have the whole thing in about two pages.

Which tools. Name them, with the account tier that counts, because the consumer version and the business version of the same product often carry opposite data terms. Say what an employee does when they want something not on the list: a two-day answer from a named person, not a refusal, because a refusal produces shadow use.

What data. Tiers work where blanket rules fail. Open: published material, marketing copy, public documents; any approved tool. Internal: drafts, plans, unpublished numbers; approved business-tier tools only. Restricted: customer personal data, anything under a confidentiality agreement, credentials, health or payment information, anything a contract says stays in a particular place; not into a general-purpose tool at all without a named approval.

Which uses. The prohibitions are shorter and more useful than the permissions. No final decision about a person (hiring, discipline, credit, termination) made on model output. No legal, medical or financial advice generated for a customer. No customer data in a personal account. No AI-generated likeness or voice of a real person without written permission. Nothing published without a human check.

What must be verified. Every factual claim, number, citation and name, by the person who sends it. Then the two lines people skip: a named owner who maintains the tool list, answers exception requests and is allowed to say yes; and a review date every six months, or whenever a new tool is adopted, because vendors change their terms without asking you.

What you can leave out: most of what enterprise governance material contains. Risk tiering, impact assessment templates, standing committees, bias-testing protocols for models you did not build. Those exist for organizations that build or deploy models at scale. At your size they produce a document nobody reads, which is worse than a short one people do. The well-known public frameworks are good structure to borrow the shape of; they are not documents to implement clause by clause.

THE CHECKLIST

- [ ] **Answer the four questions in two pages**

  Approved tools with account tiers; data tiers (open, internal, restricted); prohibited uses; the verification rule. Plain words, because staff have to follow it.

- [ ] **Name the owner**

  One person who keeps the tool list, answers requests within two days, and can say yes. Without a name the policy is a document. With one it is a process.

  > The owner is rarely the most technical person. It is the one who can read a vendor's terms and a customer contract and hold the line.

- [ ] **Set the review date**

  Twice a year, and on any new tool. Tools change their terms unilaterally; the policy has to catch up.

- [ ] **Add one sentence on what you tell customers**

  So that three account managers do not improvise three different answers when asked how the company uses AI.

- [ ] **Introduce it in one meeting, with three real examples**

  Walk the tiers using your own work, name the owner in the room, and answer the question everyone actually has: whether they are in trouble for the tool they already use. They are not, and saying so is what makes the policy real.

- [ ] **Check it against the contracts you have already signed**

  Some customer and vendor agreements already restrict AI processing of the other side's data. A policy written without reading them will contradict them (see the chapter on what you have promised).

- [ ] **Leave out the apparatus**

  No committees, tiers of model risk, or assessment templates unless you build models. A short policy people follow beats a long one they do not.

CHAPTER 4

Read the vendor's terms: the six clauses

An AI vendor's terms run to thousands of words, and six clauses decide almost everything. Read those six and you know which tier of your own policy a tool belongs in. This is worth doing even when nothing is negotiable: a click-through you cannot change still tells you whether the tool may see customer data, internal drafts, or nothing but published material.

Training rights. Almost nobody writes "we will train our models on your data." They write that the provider may use customer data to develop, maintain and improve the services. Unqualified, that is a training right, and it is the most consequential sentence in most AI agreements. Look for an explicit statement that your content is not used to train or improve models; whether that is the default or an opt-out you must switch on; whether it covers inputs, outputs or both; and whether it holds on the tier you are actually using.

Retention and deletion. How long is data kept in the ordinary course (a short period for abuse monitoring is common and usually fine), and what happens when you delete: is the record purged, or does your access end while the data lives on in logs and backups? Ask for the period as a number. "As long as necessary" is not a retention period.

Subprocessors and location. Who else touches the data, where it is processed, and whether you are told before that changes. If your own customer contracts promise data stays in a particular country, this clause either honors that or quietly breaks it. Ask separately whether people review a sample of traffic, because that may be a disclosure your customers are owed.

Output ownership and the license back. Do you own the output, and what license does the vendor keep in it? A broad license back can matter as much as a training right. Remember, too, that owning the output as against the vendor is not the same as owning it against the world: the same output may be produced for another customer, and copyright requires human authorship.

Indemnity and its conditions. Several major vendors now promise to defend you against copyright claims over their output. Read the conditions (safety features on, filters not disabled, no infringing input, unmodified output) and the cap. An indemnity capped at a year's subscription fees is not a backstop for a real claim.

Exit and change. What you can take out, in what format, and what happens to the rest. And how the terms change: many AI agreements allow amendment on notice, so the training clause you approved in the spring may not be the one in force by autumn. Diary a re-read.

Before any of the six, one prior question: what have you already promised someone else about this data? Customer contracts, confidentiality agreements and data-processing terms frequently restrict who may process the other side's information and where. A vendor's terms can be perfectly reasonable and still be incompatible with a promise you made last year to a customer who has never heard of that vendor. The vendor read is the second step. The list of what you already agreed to is the first, and for most organizations nobody has assembled it.

THE CHECKLIST

- [ ] **Before the vendor's terms, list what you already promised**

  Customer contracts, confidentiality agreements, data-processing terms, and any professional duties. This is the standard the vendor's terms have to meet.

- [ ] **Find the training clause and read it as written**

  "Develop, maintain and improve the services," unqualified, is a training right. You want an explicit exclusion, and you want to know whether it is default or a switch.

  > If the protection is a switch, name the person who throws it and confirm it in writing. A protection nobody configured is theoretical.

- [ ] **Get the retention period as a number**

  And a plain answer to what deletion actually deletes.

- [ ] **Check subprocessors and location against your own promises**

  A list, a notice commitment, and a processing location that matches what your contracts say about where data may go.

- [ ] **Read the indemnity's conditions and cap before relying on it**

  Real but often small. Know what you must do to keep it, and how much it would actually pay.

- [ ] **Diary the re-read**

  At the policy review date, and whenever the vendor notifies a change. Note whether continued use counts as acceptance.

- [ ] Write down the verdict as a tier

  Approved for open material only; approved for internal drafts; approved for restricted data with named approval. A vendor review that does not end in a tier changed nothing.

CHAPTER 5

What you own, and what you do not

This is the exposure that surprises people most. Under current United States law, a work created entirely by a machine, with no meaningful human creative contribution, is not protected by copyright at all. The Copyright Office has said so in guidance and in registration decisions, and courts have agreed when asked. A detailed prompt has not been treated as enough. A person has to have made real creative choices in the result.

Consider what that means. If your new logo was generated by a tool and adopted as is, you may hold no copyright in it, and may be unable to stop a competitor using something identical. If your website copy, product documentation or training videos were machine-generated end to end, the machine-generated part may not be protected by copyright at all. That is not the same as putting it in the public domain on purpose, and the human work around it (what you selected, how you arranged it, what you rewrote) can still be protectable on its own. Your trademark in the brand name is separate and still works; the copyright in the artwork may simply not exist.

The mirror-image problem is that output can reproduce material from the model's training data. If a tool hands you three paragraphs that closely track a copyrighted article and you publish them, you published them. The vendor's terms are careful to leave that with you, subject to whatever indemnity you read in the last chapter.

The lawsuits against the model builders continue, and the early decisions have not all pointed the same way: some judges have found that training a model on lawfully obtained works can be a fair use, others have rejected that argument on different facts, and how the copies were acquired has mattered a great deal. The honest summary is that the law of AI training is still being written. A business that uses these tools should plan for uncertainty rather than assume the vendor has it covered.

THE CHECKLIST

- [ ] **Put a person's real creative work into anything you would want to enforce**

  Brand assets, core content, product design. Use the tool for drafts and options; have a human make the choices that matter, and keep a record that they did.

  > A logo adopted straight from a generator may belong to nobody. That is fine for an internal slide. It is not fine for the mark on your storefront.

- [ ] **Register the copyright in what matters, once a human has made it theirs**

  Registration is what lets you enforce it. The application will ask about AI-generated material, and the answer has to be accurate.

- [ ] **Treat output as unverified for originality as well as accuracy**

  Before publishing anything substantial, check that it does not closely track an existing work. The person who publishes it is the one who answers for it.

- [ ] **Check your own content licenses for AI restrictions**

  Research reports, stock libraries, third-party data. Some licenses now forbid use with AI systems; silence is not permission.

- [ ] **Do not plan around the outcome of the training lawsuits**

  They may change what vendors can offer and what indemnities cover. Your policy should work whichever way they come out.

CHAPTER 6

What you have promised

Organizations talk about AI in four directions, and each is a place a statement can be untrue. To customers, in contracts and marketing. To investors, in decks and disclosures. To job candidates and employees. And to regulators, when they ask.

Look first at your standard customer contract, the one your company signs as the vendor. It probably says the work is original, that you own it, that it does not infringe anyone's rights, and that you will defend the customer if someone claims otherwise. That language was drafted for a world in which people made the deliverables. If a meaningful share of what you deliver is now machine-generated, each promise deserves a second look: "original" may not be accurate, "we own it" may not be true for the reasons in the last chapter, and the indemnity now covers a risk you did not price. Sophisticated customers have noticed. Requests for proposals increasingly ask vendors to disclose their use of AI, and some contracts now forbid it for certain work or require written consent.

Then the marketing. Regulators have brought cases for what they call AI washing: claiming a product uses AI when it does not, or claiming capabilities it does not have. "AI-powered" is a claim, and it will be read as one.

Then the people you hire. In New York City, an employer that uses an automated tool to screen or rank candidates must have that tool independently audited for bias, publish the results, and tell candidates in advance that it is being used. Other places have similar rules arriving. Recruiting software you licensed may have switched such a feature on without anyone in the company deciding to use it.

The fix in every direction is honesty in writing. Describe how the work is produced. Where AI tools are used, say so, and say what a human does before delivery. Adjust the warranties to promise what you can deliver: that the work does not knowingly infringe, that it has been reviewed by a person, that the customer owns whatever rights you have in it. Customers do not mind that you use these tools. They mind finding out later.

THE CHECKLIST

- [ ] **Reread your standard customer contract with "is this still true?" beside every warranty**

  Originality, ownership, non-infringement, indemnity. Where the answer is no, fix the language before the next renewal, not after the first claim.

- [ ] **Answer the AI question in your proposals and contracts the same way every time**

  One paragraph: which kinds of tools are used, under what data protections, what a human reviews, and that the customer may ask for a matter to be handled without them.

  > Customers rarely object to disclosed use. They object to discovering it. The paragraph is the cheapest fix in this handbook.

- [ ] **Audit your marketing for AI claims**

  Every "AI-powered," "intelligent," "learns" and "predicts" is a factual claim. Keep the ones you can prove.

- [ ] **Check whether any tool screens, ranks or scores job candidates**

  If it substantially assists or replaces a hiring or promotion decision and you employ in New York City, the law calls it an automated employment decision tool. A bias audit by an independent auditor within the past year, a summary of its results made publicly available, and notice to candidates at least ten business days before you use it are legal requirements, not best practice. The notice covers the tool and the qualifications it assesses, and candidates may ask what data it uses and how long you keep it. Screening or ranking is only one way a tool can be covered, so read the definition against your own process. Elsewhere, assume similar rules are coming.

- [ ] **Check what your investor materials say**

  Claims about AI capability or AI-driven growth in a deck or a filing are held to the same standard as any other financial statement.

- [ ] **Keep a copy of every AI statement you make to anyone**

  So that when a customer, a regulator or a board member asks what you said, there is one answer.

CHAPTER 7

Decisions about people

The uses of AI with the most law attached are the ones that decide something about a human being: who gets interviewed, who gets credit, who gets a warning, whose claim is paid, whose account is closed. Tools that do these things are useful and increasingly common, and the rules around them are stricter than around any other use.

The core principle is old law applied to new tools. Decisions that affect people's employment, credit, housing, insurance or access to services must not discriminate on protected grounds, and the fact that a model made the decision is no defense; the organization that deployed it is responsible for its outcomes. On top of that general principle, specific rules are arriving: New York City's requirement of an independent bias audit and advance notice for automated hiring tools is the clearest example, and others follow the same pattern of audit, notice and a human path.

The practical rule for an organization that does not build these systems is simple. No final decision about a person is made on model output alone. A tool may sort, flag, suggest or draft; a person decides, can see why the tool suggested what it did, and can overrule it. That rule is both what the emerging law requires and what protects the organization when a decision is challenged, because it means someone can explain it.

Two further habits matter. Keep records of what the tool recommended and what the human decided, so that patterns can be checked. And give people a way to ask for a human review of a decision that affected them, because that is what several of the new rules require and what a fair organization would offer anyway.

THE CHECKLIST

- [ ] **Write the rule down: no final decision about a person on model output alone**

  Hiring, discipline, termination, credit, eligibility, pricing that varies by person. A human decides, can see the reasons, and can overrule.

  > Rubber-stamping is not human review. If the person approving cannot say why, the tool is deciding.

- [ ] **Inventory every tool that scores, ranks or screens people**

  Including features inside recruiting, customer-service and finance software that vendors switched on. This list goes to the board.

- [ ] **If you use an automated hiring tool for New York City roles, get the bias audit and give the notice**

  Independent audit, results published, candidates told in advance. Ask the vendor for their audit; the obligation is yours regardless.

- [ ] **Keep the record of recommendation and decision**

  What the tool suggested, what the person decided, and why they differed when they did. It is how you notice a pattern before a regulator does.

- [ ] **Offer a human path**

  A way for a candidate, customer or employee to ask that a person look again. Several new rules require it; all of them reward it.

- [ ] **Ask the vendor the four questions before you rely on a decision tool**

  What data it was trained on, whether it has been tested for disparate outcomes, whether it can explain a given output, and who is responsible if it is wrong.

CHAPTER 8

Confidential, privileged, and the chatbot

Owners now take questions to a chatbot before they take them to anyone, including legal questions about the business. It is a useful habit and a dangerous one, and the difference is knowing what the tool cannot do.

Nothing you tell it is privileged. What you tell your lawyer, in confidence, to get legal advice, cannot be forced out of either of you. Nothing like that attaches to a chatbot. If a dispute ends in litigation, the chat history can be demanded. Worse, the question itself becomes evidence: type "is it legal to do X," then do X, and you have created a dated record showing you wondered whether X was legal before you did it. Ask your lawyer the same question and the exchange is protected.

Your confidential information may leave the building. Depending on the product and the settings, what you paste may be kept, seen by people, or used to train the next model. The term sheet, the customer list, the unsent letter: once they sit in someone else's system under someone else's terms, you will struggle to say you took reasonable steps to keep them secret, and reasonable steps are what trade-secret protection depends on. If the material belonged to a customer under a confidentiality agreement, you may have breached it the moment you pressed enter.

And the tool does not know you. It answers in general; your problem is specific. It does not know New York's rule that a corporation's board can act without a meeting only if every director signs, or the clause in your lease, or the fact you did not think to mention. It does not know when the law changed. It cannot check what it was not told, and it will not ask. When it is wrong, it is wrong in a confident voice, and nobody is accountable.

Used well, it is a very good way to learn the vocabulary of a problem before you call anyone, so the call is shorter and cheaper. It is good at summarizing a document you already have and at helping you write down your questions. Used that way it makes you a better client. Used as counsel, it is the most expensive free thing your business will touch.

THE CHECKLIST

- [ ] **Treat every chatbot conversation as discoverable**

  Because it may be. Do not ask it whether something you plan to do is legal, in those words, about your actual situation. Ask a lawyer, where the question is protected.

> The question is the exhibit. This is the one owners do not see coming.

- [ ] **Keep restricted material out of general-purpose tools**

  Customer data, anything under a confidentiality agreement, trade secrets, unsent negotiating positions. Your own policy's restricted tier exists for exactly this.

- [ ] **Use it to prepare, not to decide**

  Learn the terms, summarize the document you already have, draft your list of questions. Then take the questions to a person who is responsible for the answers.

- [ ] **Verify anything it tells you about the law, or about New York**

  Rules are specific and they change. A general answer in a confident voice is the tool's failure mode, and the state-specific rule is where it fails most.

- [ ] **Remember it advises both sides**

  It owes you no loyalty and runs no conflict check. It will help the other party to your negotiation the same afternoon and remember neither of you.

CHAPTER 9

If you are a professional firm

Law firms, accounting firms, medical practices and anyone else who holds client information under a professional duty face everything in this handbook plus a layer of rules specific to their profession. For lawyers, the clearest statement is the American Bar Association's formal opinion on generative AI, issued in 2024, and the guidance from New York's bar associations that followed. Neither created new duties. The opinion names seven duties a lawyer already owes and works out what each one means once a generative tool is in the workflow.

Competence: a reasonable understanding of the specific tool's capabilities and limits, not technical mastery, and no uncritical reliance on its output. Confidentiality: before client information goes into a tool, an assessment of the risk it will be disclosed, and, for a tool that learns from what users enter, the client's informed consent first, which a boilerplate waiver in an engagement letter does not supply. Communication: no standing duty to announce AI use, but disclosure when the client asks how the work was done, when the engagement terms or the client's own rules for its outside lawyers require it, when client information is entered, when the use affects the fee, and when the output will shape a significant decision in the matter; the engagement letter is the natural place. Meritorious claims and contentions: a position you would not have signed after reading the authority yourself is not one a tool can make arguable. Candor: the filings that cited invented cases were lawyers signing what they had not checked; New York's courts have since adopted a statewide policy (Part 161). It says two things: using AI to prepare a submission should not be prohibited, and nobody should have to tell the court they used it. The review requirement people quote is not in that policy. It sits in a model rule at Appendix A, and section 161.4 leaves each court free to adopt it or not. Where a court has adopted it, the filer must carefully review the paper and independently ensure it contains no fabricated or fictitious cases, statutes or other material, and signing certifies that the review happened. Where a court has not, the older duties still bite, because signing a paper already certifies it carries no false statement of fact and no frivolous argument. Read the part rules of the judge you are actually filing in front of. Supervision: managing lawyers must establish clear policies on permissible use and make reasonable efforts to see that lawyers and staff follow them, which is what a malpractice carrier is testing when it asks whether the firm has a written AI policy. Fees: bill the time actually spent; if a tool turns four hours into one, the client is billed one, and the hours spent learning a tool you will use across the practice are the firm's, not the client's. The exception runs the other way: where a client asks you to use a particular tool you do not know, the time to learn that one can be billable.

In practice the two that take real work are the consent question for self-learning tools, which means someone has to read the terms of every tool and know which ones learn, and the supervision requirement, which means the policy is not optional. What the opinion does not require is worth stating too, because firms overcorrect: no blanket disclosure to courts or opponents, no consent for every use, no particular vendor or committee, and no requirement to stop.

For a six-lawyer firm the work looks like this. List the tools (three, it turns out, one on an associate's personal account, which stops today). Read the terms of each; one learns from inputs on the consumer tier and not on the business tier, so upgrade or stop putting matter information into it. Add a paragraph to the engagement letter. Where a self-learning tool will see matter information, have the consent conversation and note it in the file. Write two pages of policy. Fix the billing narrative. Train everyone once and put the date in the file, because that date is what the carrier asks for. The first, second and sixth items are the work; the rest is drafting.

THE CHECKLIST

- [ ] **Map your tools against the seven duties**

  Competence, confidentiality, communication, advancing only meritorious claims and contentions, candor toward the tribunal, supervision, and reasonable fees. The consent question and the written policy are usually where the work is.

- [ ] **Know which of your tools learn from what you enter**

  Read the terms, not the marketing page, and check the tier you actually use. A self-learning tool needs the client's informed consent before matter information goes in.

  > The same product can learn on the consumer tier and not on the business tier. An associate's personal account is the most common way a firm is out of compliance without knowing it.

- [ ] **Put one paragraph in the engagement letter**

  How the firm uses these tools, what goes into them and what does not, that a lawyer reviews all work, and that the client may ask for their matter to be handled without them.

- [ ] **Verify every citation, quotation and factual claim before it leaves the firm**

  New York's court rule permits the tools and puts the responsibility for what is filed squarely on the filer. Check each court's standing orders as well; some require certifications.

- [ ] **Write the policy and train everyone, once a year**

  Two to four pages; the four questions from the earlier chapter. The training date goes in a file, because the insurer will ask for it.

- [ ] **Bill the time actually spent**

  If the tool saves an hour, the invoice reflects an hour saved. Decide whether a tool's cost is overhead or a disbursement the way you would for any other service.

CHAPTER 10

The board's five questions, and keeping it going

A board is not asked to understand how the technology works. It is asked to know where the organization is exposed, to satisfy itself that someone competent is managing that exposure, and to write down that it did. That is the standard for every other risk a board oversees, and it is the standard for this one. Directors of New York corporations and nonprofits owe a duty of care that the courts have long read to include oversight: a board cannot assume management is handling a significant risk; it has to make a reasonable effort to be informed.

Five questions, asked in one meeting and answered from the earlier chapters, meet that standard. Where are we using it? (The inventory: the three-column table.) What goes in and what comes out? (The policy's tiers and the verification rule; the vendor reads.) What have we told people? (Contracts, marketing, candidates, investors.) Who is accountable? (The named owner, and how this reaches the board.) What happens when it goes wrong? (The incident plan, and whether the insurance covers it.)

Record the answers. The minutes should show that the board discussed AI use and risk, what management reported, what the board asked and what follow-up was agreed. Then set a date to return to it, once a year for most organizations and more often for one that builds AI into its products. This is not paper for its own sake. When something goes wrong, the difference between a board that can show a record of informed attention and one that cannot is the difference between a difficult week and a claim against the directors personally.

Something will go wrong eventually: a tool leaks confidential information, an automated decision is challenged, an employee publishes something the tool invented, a vendor changes its terms overnight. The organizations that handle it well are the ones that decided in advance who would find out, who could pause a tool, and how the customer or the regulator would be told. That is the last item on the checklist, and it is the one that turns a policy into a practice.

THE CHECKLIST

- [ ] **Put the five questions on the next board agenda, with the inventory sent in advance**

  Where we use it; what goes in and out; what we have told people; who is accountable; what happens when it goes wrong. One hour of the board's time moves the organization from unmanaged to managed.

- [ ] **Name the executive who owns AI risk across the organization**

  Not a chief AI officer; a person with authority who can describe the policy, the tools, the training and how new uses are evaluated.

- [ ] **Ask whether the incident plan and the insurance cover an AI failure**

  Cyber and errors-and-omissions policies vary and some now exclude it. Ask who has authority to pause a tool and whether they know they have it.

  > The answer "we are not sure" is a finding, and a fixable one. The answer that costs is the one nobody asked for.

- [ ] **Minute it**

  What was reported, what was asked, what was agreed, when the board returns to it. The record of informed attention is the board's protection.

- [ ] **Review everything twice a year**

  The inventory, the policy, the vendor terms, the customer paragraph, the training. The tools will have changed; so will the organization's uses of them.

- [ ] **Keep the tone right**

  This is not a program to stop people using useful tools. It is the set of boundaries that lets them use the tools with confidence, because everyone knows where the lines are.

That is the whole of it. An inventory, two pages of rules, a vendor read, a look at what you own and what you promised, a human deciding anything about a person, restricted material kept out of general tools, and a board that asked and wrote it down. None of it slows anyone. All of it is what a careful organization would want in place before the question is asked by someone outside it.

If you would like help with any of it, the conversation usually begins with the ten-minute question from the second chapter: what is your team already using? The list is where the work starts.

This is general information about how these instruments and obligations usually work, not legal advice about your situation, and reading it does not make you a client. Deadlines and requirements turn on facts particular to you: your fiscal year, your state, the agreement actually in front of you. Confirm your own before you rely on any of it.

Clemenza Law Group, New York, NY. (212) 365-4875.

The current edition of this handbook, and the essays behind it, are at clemenzalaw.com/reference/ai-at-work.