

Stand Behind the Result

A Handbook for Small Legal Teams Using AI

Your team is already using AI. It does not have a standard for it yet.

ProvisoIQ Editorial, 2026

Table of Contents

Stand Behind the Result

A Handbook for Small Legal Teams Using AI

ProvisoIQ Editorial

First edition, 2026

(c) 2026 Proviso Inc. All rights reserved.

Information and tools only, not legal advice. Proviso Inc. is not a law firm. Reading this book creates no attorney-client relationship. Review every AI output before relying on it, and apply your own professional judgment and the rules of the jurisdictions where you practice.

Statements about laws, guidance and AI providers were checked against their sources on the dates given in the endnotes. Rules and product terms change; check the current version before you rely on any of them.

Contents

- Introduction: Where You Are Starting From
- Choosing What to Hand to AI First
- Safe Use and Confidentiality
- Prompting for Legal Work
- Checking Output
- Contract Review with AI
- Intake and Triage
- Reporting to the Business
- Building a Team Standard
- Measuring Time Saved Honestly
- Explaining Legal to the Business
- What Comes Next: Agents

- Appendix A: A Starter Prompt Library
- Appendix B: A Sample Team Standard
- Appendix C: Glossary
- Appendix D: A Thirty-Day Plan
- Appendix E: Building a Test Set
- Appendix F: Asking for Budget
- Appendix G: For Paralegals, Legal Operations and Legal Assistants
- Appendix H: For Small-Business Owners Without In-House Counsel
- Appendix I: Questions Leaders Ask About Legal's AI Use
- Appendix J: A Sample Note on How Legal Uses AI
- Endnotes

Introduction: Where You Are Starting From

Someone has asked you to use AI. Perhaps it was the general counsel, the CFO, or a chief executive who read an article on a flight. Perhaps nobody asked, and you simply noticed that the sales team now drafts emails in seconds while your queue of agreements looks the same as it did last year. Either way, the request arrived without the things that usually come with a new way of working: time to learn, a budget, training, and someone down the hall who has already done it.

If that is your situation, this book is for you. It is written for small legal teams: the department of one, the general counsel with two lawyers and a paralegal, the legal operations lead who supports a team spread across three time zones, and the legal assistant who has been told to “find ways to use AI” without being told where the line sits. It is also written for the small-business owner who has no in-house counsel and wants to know what to check before signing.

The book assumes you are careful. That caution is a strength, and nothing here asks you to give it up. Legal work is cautious because the people who rely on it need it to hold up, and a tool that produces confident text quickly does not change that need. What the book offers is a way to keep your standards while using AI on real work, so you can stand behind the result.

What this book is and is not

This is a practical handbook. Each chapter takes one part of the work, explains how AI can help with it, names the errors AI tools tend to make on that task, and ends with a short checklist you can use the same day. The examples use invented documents and invented companies, because nothing in a handbook should depend on anyone's confidential material.

It is not legal advice, and it does not tell you what the law requires in your jurisdiction. Where the book mentions professional guidance or regulation, it says where the point comes from and when it was checked, and the endnotes give the sources. Rules about AI are changing quickly, so treat every legal reference as a pointer to check, not a conclusion to rely on.

It is also not a product review. The book does not rank AI models or vendors, and nobody paid for any mention. The prompts and methods are written to work in the general AI tools most

organizations already license, and where a tool behaves differently, the book says so in general terms and asks you to test it yourself.

The method in one paragraph

Pick one task that repeats. Write down what a good result looks like before you start. Give the AI the material it needs and nothing more, in a structured prompt.

Read the output against your own standard, using a short list of the errors that task tends to produce. Record what worked, what failed and how long it took, and change one thing at a time. When the method works for you, write it down so the rest of the team can use it the same way.

That is the whole approach. The chapters that follow fill in each step with the detail a small team needs.

How to read it

You can read the book in order, and the early chapters build the habits the later ones depend on. If you are short of time, start with Chapter 2 to choose a first task, read Chapter 3 before you paste anything into a tool, and then go straight to the chapter for your task. Chapter 5, on checking output, is the one to come back to.

If you lead a team, Chapter 9 covers turning individual habits into a shared standard, and Chapters 8 and 11 cover the conversations with the business that usually decide whether AI use in legal is supported or questioned. Chapter 10 covers measuring time saved without overstating it, which matters if you will ever ask for budget.

A word about time

Time is the reason most people pick up AI tools, and it is a fair reason. For many legal tasks, the slowest part is starting: the blank page, the first read of a long agreement, the first draft of a summary nobody will thank you for. AI can take much of that starting effort away, and a first pass you can review is often quicker to fix than a page you have to fill.

The time saving is real only if the review holds. A first draft that looks finished and hides one wrong cap amount has not saved you anything. That is why this book pairs every time-saving method with a check, and why Chapter 10 asks you to measure your own results rather than rely on anyone's general claim.

A word about tone

The book is written by people who do this work, for people who do this work. It does not assume you are behind, and it does not assume you are an expert. Where something is hard, it says so; where something is simpler than it looks, it says that too.

If you are starting from nothing, that is a perfectly good place to start. One task, done carefully, is enough to begin.

Checklist: before you start

- Write down one recurring task that takes more of your week than it should.

- Find out which AI tools your organization has approved, and on which plan.
- Find your organization's AI use policy, if there is one, and read it once.
- Set aside 30 minutes for your first test, on invented or public material.
- Decide where you will record what you learn, even if it is a single page.

Choosing What to Hand to AI First

The first task you choose shapes everything that follows. Choose well, and you get a quick, honest result that teaches you how the tool behaves. Choose badly, and you spend an afternoon on a task the tool handles poorly, conclude that AI does not work for legal, and return to the old way with a story that will be hard to change later. This chapter is about choosing well.

What makes a good first task

A good first task has five features. It repeats, so what you learn pays off more than once. It has a clear standard, so you can tell a good result from a bad one.

It is low in consequence if the first attempt goes wrong, because you will review it anyway. It can be done on material that is safe to use. And its slowest part is the start, not the judgment.

Think of the work that sits in your queue every week. The NDA that arrives with the same three problems each time. The summary of a services agreement for a business owner who will read only the first paragraph.

The monthly report to leadership that takes a morning to assemble. The intake email that needs sorting before anyone can act. These are good candidates because you already know what a good result looks like.

What makes a poor first task

Avoid tasks that turn on judgment the tool cannot have. A negotiation strategy that depends on what you know about the counterparty's last three deals is a poor first task, because the tool lacks that context and you will spend more time explaining it than doing the work. So is anything that requires legal research with citations, at least until you have a reliable way to check every source.

Avoid tasks where the consequence of an unnoticed error is high and the error is hard to see. A clause that looks right but has one number wrong, in a document going out for signature the same day, is the worst kind of first experiment. Start where you have time to check properly.

Avoid tasks that require pasting material you are not allowed to share. Chapter 3 covers confidentiality in detail, and the short version is that your first tests should use invented, public or properly cleared material.

A simple sorting method

Take a page and list ten tasks from your last two weeks. For each one, answer four questions with a yes or no.

1. Does it repeat at least monthly?
2. Can I describe a good result in three or four sentences?
3. Would I review the result anyway before it goes anywhere?
4. Can I test it on material that is safe to use?

Any task with four yes answers is a candidate. Among the candidates, pick the one where starting is the slowest part. That is your first task.

If no task scores four, you have learned something useful: either your work is unusually judgment-heavy, or the material you handle needs a cleared tool before you can test. In the second case, Chapter 3 will help you ask the right questions of whoever controls the tools.

Common first tasks that work

Several tasks tend to work well for small legal teams, and the later chapters cover each in detail.

A first-pass summary of an inbound agreement. The tool reads the agreement and produces a structured summary: parties, term, payment, liability, termination, governing law and anything unusual. You check the summary against the document. Chapter 6 covers this.

An issues list against your playbook. If you have standard positions, the tool can compare an agreement to them and list deviations. The quality depends on how clearly your positions are written, which is a useful discipline in itself.

A plain-language explanation for a business colleague. You have the analysis; the tool helps you say it in a way a sales manager will read. Chapter 11 covers this.

Sorting incoming requests. The tool reads requests and suggests a category, urgency and the information missing. Chapter 7 covers this.

A first draft of a recurring report. You supply the facts; the tool drafts the structure and wording. Chapter 8 covers this.

Set the standard before you start

Before your first test, write down what a good result looks like. For an agreement summary, it might be: every field filled from the document; nothing stated that the document does not say; unusual clauses flagged with their section numbers; under one page. For a business explanation, it might be: the decision in the first sentence; no legal terms without a plain explanation; the action required and by when.

Writing the standard first does two things. It gives you a fair test, because you will judge the output against what you needed rather than against whatever the tool happened to produce. And it becomes the first part of your prompt, as Chapter 4 explains.

Run the first test

Use invented or public material. Run the task once the way you normally would, and note roughly how long it took. Then run it with the AI tool, using a structured prompt, and note how

long it took including your review. Compare the two results against your written standard, not against each other.

Expect the first attempt to be imperfect. The point of the first test is to learn how the tool fails on this task, not to prove that it never fails. Note each error and what would have caught it. Those notes become the “what to check” list for the task.

Then change one thing, such as a clearer instruction or a better structure for the output, and run it again. Two or three rounds usually show whether the task is worth adopting.

When to stop

Stop and choose another task if, after three rounds, the review takes longer than doing the work yourself, or the errors are the kind you would struggle to spot. That is not a failure. It tells you where the tool’s limits sit for your work, and that is worth knowing.

Stop and celebrate, quietly, if the first pass is reliably ready to review and your check catches the errors that remain. You have a method. The next step is to write it down so you can repeat it.

A worked example: the department of one

Consider an invented company, Harbor Lane Software, with one in-house lawyer and a part-time legal assistant. Over two weeks, the lawyer lists the tasks that filled her time: fourteen NDAs from prospective customers and partners, six vendor agreements for the operations team, two employment questions, one customer dispute, a board update, a privacy question about a new marketing tool and a steady trickle of “can you look at this?” emails.

She scores each task on the four questions. The NDAs score four out of four: they repeat weekly, she can describe a good result (every term mapped against her standard positions, departures listed, a reply drafted), she reviews every one anyway and she can test on invented NDAs. The vendor agreements score three, because she is not yet sure the company’s AI tool is cleared for vendor pricing. The dispute and the employment questions score low, because they turn on judgment and contain sensitive facts.

She picks the NDAs. Her written standard is four sentences: every clause in the standard NDA topics is identified; each departure from her positions is listed with its section number and quoted wording; nothing is stated that the NDA does not say; and a short reply to the counterparty is drafted for each departure she will raise.

She builds five invented NDAs, deliberately varied: one mutual and clean, one one-way with a perpetual term, one with a broad non-solicitation clause, one with a residuals clause, and one badly drafted with inconsistent defined terms. She reviews each by hand first and times herself. Then she runs each through the tool with a structured prompt and her positions, and times the whole process including her check.

The first round produces useful first passes on the clean and perpetual-term NDAs. It misses the residuals clause entirely, describing the agreement as standard, and it reads the badly drafted NDA’s inconsistent definitions as if they were consistent. She adds two lines to her prompt (“List any clause permitting use of information retained in memory” and “If a defined

term is used inconsistently, say so and quote each use”) and two items to her “what to check” list. The second round catches both.

After three rounds, she has a method she trusts for NDAs, a “what to check” list of six items, and honest timings for both ways of working. She has also learned something about the tool’s behavior that will help with the vendor agreements once the confidentiality question is settled.

Questions people ask

What if my work is all judgment? Most legal work mixes judgment with routine. Even a complex negotiation has a first read, a summary and a follow-up note. Look for the routine parts around the judgment, and start there.

Should I start with something impressive, to persuade skeptics? It is tempting, but a modest first task done well is more persuasive than an impressive one done badly. Skeptics are persuaded by results they can check, and a reliable NDA first pass is easier to check than a novel use.

What if my organization has no approved AI tool? You can still prepare. Build your written standards, playbook positions and invented test documents now. When a tool is approved, you will be ready to test quickly, and your preparation will help whoever is choosing the tool understand what legal needs.

How long should the first test take? Plan for two or three sessions of an hour or less. If the method is not taking shape after that, choose another task. Long first experiments tend to produce fatigue rather than learning.

What if the tool does better than I expected on a task I thought was too hard? Treat the result with care. A strong result on a few test documents is encouraging, but judgment-heavy tasks can fail in ways a small test does not show. Expand testing before relying on it, and keep the review standard high.

Checklist: choosing a first task

- List ten tasks from the last two weeks.
- Score each on the four questions: repeats, clear standard, reviewed anyway, safe material.
- Pick the candidate where starting is the slowest part.
- Write down what a good result looks like in three or four sentences.
- Test on invented or public material, and time both the old way and the new way.
- Record each error and what caught it.
- Decide after three rounds whether to adopt, adjust or move on.

Safe Use and Confidentiality

Before you paste anything into an AI tool, you need to know where it goes. That question sounds technical, and the answer often sits in terms of service that few people read. It is also

the question most likely to decide whether your use of AI is defensible. This chapter sets out what to find out, who to ask, and how to work safely while the answers come in.

Why this comes first

Lawyers have a duty to protect client information, and in-house lawyers usually owe a similar duty to their organization under employment terms and policy. Legal staff who are not lawyers work under the same expectations through their employer and through the lawyers who supervise them. Pasting a document into a tool is a disclosure to whoever runs that tool, and the question is whether that disclosure is permitted and protected.

Professional guidance now addresses this directly. The American Bar Association's Formal Opinion 512, issued on July 29, 2024, discusses how the duties of competence, confidentiality, communication and reasonable fees apply when lawyers use generative AI tools.¹ The State Bar of California's practical guidance on generative AI, first issued in November 2023 and since updated, says that a lawyer must make reasonable efforts to protect client confidences when using these tools, and that this requires understanding how a product collects, uses, stores and discloses information, beyond relying on general marketing assurances.² Other jurisdictions have issued their own guidance, so check the rules that apply to you.

The five questions to answer for any tool

For each AI tool you might use, find the answers to five questions. Write them down with the date you checked, because terms change.

1. **Which plan is this?** Consumer, team, business and enterprise plans of the same product often have different terms. The free version you use at home may treat your data very differently from the business version your company licenses.
2. **Is input used to train the provider's models?** Find the current setting and the current terms. Some plans exclude business data from training by default; some let users opt out; some do not.
3. **How long is input kept, and can you change that?** Retention periods and deletion controls vary by plan and sometimes by feature.

¹ American Bar Association, Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512, "Generative Artificial Intelligence Tools" (July 29, 2024), addressing competence (Model Rule 1.1), confidentiality (Model Rule 1.6), communication (Model Rule 1.4) and fees (Model Rule 1.5), among other duties. ABA announcement: <https://www.americanbar.org/news/abanews/aba-news-archives/2024/07/aba-issues-guidance-lawyers-use-ai-tools/>

² State Bar of California, Standing Committee on Professional Responsibility and Conduct, "Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law" (originally issued November 2023; current version on the State Bar website), including the statement that a lawyer must make reasonable efforts to protect client confidences when using generative AI. <https://www.calbar.ca.gov/sites/default/files/portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf>

4. **Who at the provider can see it, and when?** Terms usually allow limited human review for safety or abuse monitoring. Know what the terms say.
5. **Is there a signed agreement?** A negotiated agreement with confidentiality and data processing terms gives you something firmer than a click-through.

You may not be the person who can answer all five. In many organizations, IT or procurement holds the contract and can tell you which plan you are on. Ask them. A short, specific question (“Are we on the business plan for this tool, and is training on our data turned off?”) usually gets a faster answer than a general one.

Sorting what you can paste

Once you know the answers, sort information into three groups.

Safe to use in any approved tool. Public material, invented examples, your own templates and playbooks with no client details, and general questions. This is where your first tests belong.

Use only in a tool cleared for it. Company confidential material, draft agreements with counterparties named, internal reports. Use these only in a tool your organization has approved for confidential information, on the plan that approval covers.

Do not paste. Material you are not permitted to disclose to any third party, privileged communications unless your organization has specifically cleared the tool for them, personal data beyond what the task needs, and anything covered by a confidentiality obligation that forbids sharing with service providers. When in doubt, it belongs here until someone with authority says otherwise.

Reducing what you share

Most tasks need less information than you think. A clause review rarely needs the counterparty’s name. A summary of a services agreement rarely needs the bank details in the payment schedule. Before pasting, remove what the task does not need: replace names with roles (“Customer”, “Vendor”), remove personal data, and cut sections that are not relevant.

This habit has a second benefit. Shorter, cleaner input usually produces better output, because the tool is not distracted by material that does not matter.

Privilege

Privilege is the area where guidance is least settled, and the safest approach is conservative. Disclosure to a third party can risk waiver, and whether a given AI tool counts as a disclosure, or as a service provider acting for you, can depend on the terms, the jurisdiction and the facts. Until your organization has a considered position, keep privileged communications and legal advice out of general AI tools, and treat any position you do take as one to revisit as courts and regulators address it.

Telling the people who need to know

Some guidance suggests that lawyers consider whether to tell clients about AI use, particularly where it affects the work, the risk or the fee. For in-house teams, the “client” is usually the organization, and the practical step is to make sure your use fits the organization’s AI policy and that your manager knows how you are using the tools. If you work in a firm, check the guidance in your jurisdiction and your engagement terms.

Working safely while you wait

You do not have to wait for every answer before you start learning. Use invented documents and public material to build your prompts and your “what to check” lists. When a cleared tool becomes available, you will be ready to use it well on real work, and you will have a record that shows you approached it carefully.

A worked example: clearing a tool

The lawyer at the invented Harbor Lane Software wants to use the company’s AI tool on vendor agreements. She writes to the IT manager with five short questions: which plan the company is on; whether input is used for training on that plan; how long input is kept and whether that can be shortened; what the terms say about provider access to input; and whether there is a signed agreement with data protection terms.

The answers come back in two days. The company is on a business plan. The plan’s terms exclude business input from training by default, and the setting is confirmed off in the admin console.

Retention follows the provider’s standard business terms, which the IT manager links. Provider access is limited to safety and abuse review under the terms. There is a signed order form that incorporates the provider’s business terms and data processing addendum.

She records the answers with the date and the links, and she reads the linked terms herself to confirm the key points. She then writes a short note for her own file: the tool is cleared for company confidential material, including draft vendor agreements with counterparty names; privileged communications stay out until she has considered the question further; personal data is removed unless a task requires it.

She also notes what would change the position: a move to a different plan, a change in the provider’s terms, or a new feature that sends data to a different service. She sets a reminder to recheck the terms every six months, and sooner if the provider announces a change.

A second example: the shared document

A legal assistant at another invented company is asked to summarize a customer contract. The contract contains the customer’s pricing, personal contact details for the customer’s staff and a schedule describing a security incident the customer had the previous year. The assistant’s company has an approved AI tool, but the assistant is unsure whether the incident schedule should be shared.

Applying the sorting in this chapter, the assistant removes the personal contact details, which the summary does not need, and replaces the customer's name with "Customer". The incident schedule is a closer question. The contract's confidentiality clause covers information about the customer's security, and the assistant cannot tell whether the approved tool falls within the clause's permitted disclosures.

The assistant summarizes the main agreement with the tool, leaves the schedule out and notes in the summary that the schedule was reviewed separately. Then the assistant asks the supervising lawyer whether the tool may be used for that kind of material in future. That is the right outcome: the routine part of the task gets done, and the judgment question goes to the person responsible for it.

Questions people ask

Is it ever safe to use a free consumer AI tool for work? For invented material, public information and general questions, a consumer tool may be acceptable under your organization's policy. For anything confidential, the terms of consumer plans are often less protective than business plans. Check your policy and the plan's terms before using any tool for work material.

Does removing names make a document safe to paste? It reduces the risk but may not remove it. A contract with names removed can still reveal the deal, the price and the parties to anyone who knows the context. Removing names is a good habit, not a substitute for using a cleared tool.

What about AI features built into tools we already use, such as email or document software? These features may run under the same terms as the host software or under separate AI terms. Ask the same five questions about each feature. Do not assume a feature is covered by the agreement for the product it sits in.

Who decides whether a tool is cleared? Usually a combination of IT, security, procurement and legal. In a small organization, the legal team may be the one deciding, which makes a written record of the decision and its basis especially useful.

Should I tell my manager how I am using AI? Yes. A short note of the tool, the tasks and the checks you use protects you and helps your manager answer questions from others. It also makes it easier to build a team standard later.

Checklist: before you paste

- Name the tool and the plan you are using.
- Check, and date, the training, retention and review terms for that plan.
- Confirm the tool is approved for the kind of information you want to use.
- Remove names, personal data and anything the task does not need.
- Keep privileged material out unless your organization has cleared the tool for it.
- Record what you checked and when, in one place.

Prompting for Legal Work

A prompt is an instruction, and lawyers already know a great deal about writing instructions. You draft scopes of work, you brief outside counsel, you write instructions for paralegals. The same habits make good prompts: say who is doing the work, what the work is, what material to use, what the result should look like and what to avoid. This chapter turns those habits into a structure you can reuse.

The five parts of a working prompt

Most effective legal prompts have five parts. You do not need headings for each, but you should be able to point to each one.

1. Role and context. Tell the tool the setting. “You are assisting an in-house lawyer at a mid-sized software company, reviewing an inbound NDA from a prospective customer.” This shapes vocabulary, assumptions and emphasis.

2. The task. Say exactly what you want. “Identify every clause that departs from the positions listed below, and for each one explain the departure in one sentence.” Vague tasks (“review this NDA”) produce vague output.

3. The material. Provide the document or text, clearly separated from your instructions. Many people put the document between markers such as “DOCUMENT START” and “DOCUMENT END” so the tool does not confuse the document’s words with yours.

4. The standard and format. Describe the result. “Return a table with four columns: section number, clause topic, departure, suggested position. Quote the clause wording exactly. If a clause is missing, say so in its own row.” The more precisely you describe the format, the easier the output is to check.

5. Limits. Say what the tool should not do. “Do not assume facts not in the document. Do not cite case law. If you are unsure whether something is a departure, list it and say why you are unsure.”

Why the limits matter most

AI tools are built to be helpful, and an unguided tool will often fill gaps with plausible material. In legal work, plausible material that is not in the source is the most dangerous kind of error. Explicit limits reduce it.

Three limits are worth including almost every time. Tell the tool to rely only on the material provided. Tell it to quote source wording when it makes a claim about a document. And tell it to say when it is unsure, rather than guessing. These instructions do not eliminate errors, but they make the errors easier to see, because a quoted clause can be checked and a flagged uncertainty invites a look.

A worked example

Here is a weak prompt and a stronger version for the same task.

Weak: “Summarize this services agreement.”

Stronger: “You are helping an in-house lawyer prepare a one-page summary of an inbound services agreement for a business manager. Using only the agreement between the markers below, fill in these fields: parties and their roles; services; term and renewal; fees and payment terms; liability cap and what is excluded from it; indemnities; termination rights; governing law and disputes; anything unusual. For each field, give the section number and quote the key words. If a field is not addressed in the agreement, write ‘Not addressed’. Do not add commentary or advice. Keep the whole summary under 400 words.”

The stronger version takes a minute longer to write. It saves far more than a minute in checking, because every claim points to a section and quotes it.

Prompts as templates

Once a prompt works, keep it. A prompt you can reuse is a template, and templates are where the time saving compounds. Store each working prompt with a note of the task, the model you used, the date you last tested it and the errors it tends to produce.

Keep placeholders obvious, for example [AGREEMENT TYPE], [OUR SIDE] and [POSITIONS], so you remember to fill them in. A prompt that silently carries over the last deal’s details is a source of errors.

Your playbook as part of the prompt

If you have standard positions, they belong in the prompt. Write each position plainly: the topic, your standard, your fallback and the point at which you would walk away or escalate. “Liability cap: standard is fees paid in the prior 12 months; fallback is two times annual fees; escalate if uncapped or below fees paid.”

Writing positions this way is useful well beyond AI. Many small teams carry their positions in one person’s head, and putting them on paper is the first step to a team standard, as Chapter 9 explains.

Differences between tools

The major general AI tools follow structured prompts well, but they differ in how much text they accept at once, how they handle uploaded files, how they format tables and how closely they follow length limits. A prompt that works well in one tool may need small changes in another.

Test each prompt in the tool you will use, and note any changes. If your organization licenses more than one tool, keep a version of each prompt for each one rather than assuming they are interchangeable.

Long documents

Long agreements can exceed what a tool handles well in one pass, even when the tool accepts the whole file. Signs of trouble include sections that are skipped, summaries that become vaguer toward the end and clauses from early in the document that are misremembered later.

For long documents, work in parts. Ask for the summary of sections 1 to 10, then 11 to 20, and combine the results. Or ask a targeted question (“Find every clause that limits or excludes liability, and quote each one”) rather than a general one.

Follow-up questions

The first answer is not the last. If the output misses something, ask about it directly: “You did not mention the auto-renewal clause in section 12. Quote it and explain the notice period.” If an answer seems wrong, ask the tool to show the wording it relied on.

Follow-up questions are also a good check. A tool that quotes the clause accurately when asked is easier to trust than one that cannot find the wording it claimed to summarize.

A worked example: improving a prompt in three rounds

A legal operations specialist at an invented company wants a prompt that turns incoming vendor agreements into a one-page summary for budget holders. Her first prompt is short: “Summarize this vendor agreement for a non-lawyer. Focus on cost, term and risks.”

The first output is readable but uneven. It describes the fees well, mentions the term but not the auto-renewal, and lists “risks” that are general statements about vendor agreements rather than features of this one. It does not say where in the agreement anything comes from, so she has to search the document to check each point.

For the second round, she adds structure: a list of fields (vendor and service; fees and how they change; term, renewal and notice; termination; liability cap and exclusions; data and security; anything unusual), a requirement to give the section number and quote key words for each field, and a word limit. The output is now easy to check. It catches the auto-renewal. It still describes one liability cap as “reasonable”, and it misses that a price increase clause allows increases “at the vendor’s discretion”.

For the third round, she adds limits: “Do not describe any term as reasonable, standard or favorable. For any clause that allows a party to change price, scope or terms, quote the clause in full.” The third output passes her check on all five test agreements. She saves the prompt with the date, the model, and a “what to check” list that includes price-change clauses and the renewal notice window.

A template you can adapt

The structure below works for many document summaries. Replace the bracketed parts.

“You are assisting [ROLE] at [TYPE OF ORGANIZATION]. Your task is to prepare a summary of the [DOCUMENT TYPE] between the markers below for [READER], who needs to [PURPOSE]. Using only the document, complete these fields: [FIELDS]. For each field, give the section number and quote the key words. If the document does not address a field, write ‘Not addressed’. Do not describe any term as reasonable, standard, favorable or unfavorable. Do not add advice. If you are unsure how to read a clause, say so and quote it. Keep the summary under [LENGTH] words.”

Use the same structure for issues lists, comparisons and first drafts, changing the task and fields.

Questions people ask

Do I need to learn special prompt techniques? Most of what matters is clear instruction, which you already know how to write. Specialized techniques can help in some cases, but structure, examples and explicit limits do most of the work.

Should I give the tool examples of good output? Often, yes. An example of the format you want, using invented content, can improve consistency. Make sure the example does not contain real confidential information, and check that the tool does not copy details from the example into its output.

How long should a prompt be? As long as it needs to be to describe the task, the material, the format and the limits. Legal prompts that work are often several paragraphs. Length is not a problem; vagueness is.

Can I ask the tool to write the prompt for me? You can ask for a draft, and it may suggest a useful structure. Edit it so it reflects your standard, and test it like any other prompt.

What should I do when a tool update changes the output? Retest your saved prompts on your test documents, compare the results with your records, and adjust the prompt or the “what to check” list. This is one reason to keep test documents and dates with each prompt.

Checklist: writing a legal prompt

- State the role and context in one or two sentences.
- State the task precisely, with the output you need.
- Separate the material clearly from your instructions.
- Describe the format: fields, table columns, length.
- Require section numbers and quoted wording for claims about documents.
- Tell the tool to rely only on the material and to say when it is unsure.
- Save the working prompt with its task, model, test date and known errors.

Checking Output

This is the chapter to come back to. Everything else in the book makes AI output faster or easier to produce; this chapter is about whether you can rely on it. The method is simple: know the errors the task tends to produce, check for them deliberately and keep a record of what you find.

Why AI errors are different

People make errors too, but they tend to be visible. A tired reader skips a clause, and the summary has a gap. AI errors are often the opposite: the output is complete, fluent and confident, and one detail is wrong. The tone of the output gives no signal about its accuracy, so the usual cues you rely on when reviewing a colleague’s work do not apply.

This means you cannot check AI output by reading it and seeing whether it sounds right. You check it against the source, deliberately, using a list.

The errors to look for

Across legal tasks, AI errors fall into a few families. Learn them and you will spot most problems.

Invented content. The output states something the source does not say: a clause that is not there, a defined term with the wrong meaning, a citation that does not exist. This is the family people worry about most, and the limits in Chapter 4 reduce it.

Missing qualifications. The output reports the main rule and drops the exception. A liability cap is summarized without its carve-outs; a termination right is described without its notice period; an indemnity is described without its cap. In legal documents, the qualification is often the point.

Wrong numbers and dates. A cap of 12 months becomes 6; a 60-day notice period becomes 30; a renewal date is calculated from the wrong starting point. These are easy to miss because the sentence around them is correct.

Misread structure. Cross-references followed to the wrong section; a definition applied in a section where a different definition governs; a schedule treated as part of the main terms when the agreement says otherwise.

Assumed context. The tool assumes a governing law, a party's role or a standard practice that the document does not state. The output reads as if it were drawn from the document.

Overstated conclusions. The output says a clause is “standard”, “market” or “favorable” without any basis, or turns a drafting observation into legal advice.

A checking method that works

Use the same order every time.

1. Check the structure first. Does the output cover every field or section you asked for? A missing field is easier to spot before you start reading for content.

2. Check every number and date against the source. Caps, periods, amounts, dates and percentages. This takes a few minutes and catches a large share of serious errors.

3. Check every quoted passage. Does the quoted wording appear in the document, word for word? If the prompt required quotes, this check is fast.

4. Check the qualifications. For each important clause, look at the source and ask what limits, exceptions or conditions apply, then see whether the output reports them.

5. Check what the output says is missing. If the output says a clause is not addressed, search the document for it. Tools sometimes miss clauses placed in schedules or definitions.

6. Check conclusions and labels. Remove or rewrite any description the source does not support.

The “what to check” list

For each task you adopt, keep a short list of the errors it tends to produce. Start with the families above, and add the specific errors you find in testing. A good list is short enough to use every time, often five to eight items.

Put the list at the top of your prompt file or template, where you will see it every time. A list that sits in a separate folder is a list nobody uses.

How much checking is enough

The level of checking should match the consequence. A first-pass summary for your own understanding may need the structure, numbers and qualifications checks. A summary going to a business owner who will act on it needs every step. Anything going to a counterparty, a regulator or a court needs the same review you would give work you wrote yourself, because it is work you are putting your name to.

As you gain experience with a task, you will learn where the tool is reliable and where it is not. That experience can reasonably shape where you spend your checking time, but it should not replace the check.

Recording what you find

Keep a simple log: the task, the date, the tool, the errors found and what caught them. Over a few weeks, the log shows which errors recur, which prompts need improving and which tasks are reliable enough to rely on with lighter review.

The log also protects you. If anyone asks how you use AI and how you control the risk, a dated record of tests and checks is a better answer than a general assurance.

When to rewrite rather than fix

Sometimes the output is wrong in ways that make it faster to start again. If the structure is wrong, more than a few numbers are wrong or the output has misread what the document is, discard it. Improve the prompt, or do the task yourself, and note what happened in your log.

A worked example: the cap that looked right

A first-pass summary of an invented services agreement reads: “Liability: each party’s liability is capped at the fees paid in the 12 months before the claim (section 11.2).” The sentence is accurate as far as it goes, and it quotes the right section.

Following the checking method, the reviewer opens section 11 and reads it in full. Section 11.3 says the cap does not apply to breach of confidentiality, indemnification obligations or a party’s gross negligence or willful misconduct. Section 11.4 excludes indirect and consequential damages, but section 11.5 says the exclusion does not apply to data breaches caused by the

vendor. And the definition of “Fees” in section 1 excludes pass-through costs, which make up much of the contract’s value.

None of this appears in the summary. Each omission changes the picture: the vendor’s real exposure for a data breach is much larger than the summary suggests, and the cap itself is smaller than the total contract value implies. The reviewer adds the three qualifications to the summary, adds “cap carve-outs, exclusions and the definition of the cap’s measure” to the “what to check” list, and adjusts the prompt to ask for “every exception to the cap and to the exclusion of damages, quoted in full”.

The lesson is not that the tool was careless. It answered the question it was asked. The lesson is that the qualification is often the point, and the check has to look for it deliberately.

A worked example: the clause that was not there

In another invented review, a summary of a software license states: “The customer may terminate for convenience on 30 days’ notice (section 14.2).” The reviewer checks section 14.2 and finds a termination right for breach, not for convenience. There is no termination for convenience anywhere in the agreement.

When asked to quote the wording it relied on, the tool quotes section 14.2 accurately and then acknowledges that it does not provide a convenience right. The error came from the tool filling in a common clause it expected to see. Because the prompt required section numbers, the check took a minute. Without them, the error might have survived into advice to the business that the contract could be exited easily.

Questions people ask

How long should checking take? It depends on the task and the document. For a structured summary with section numbers and quotes, checking is usually much faster than doing the task from scratch. If checking regularly takes as long as doing the task, revisit the prompt or the choice of task.

Can I ask the tool to check its own work? You can ask it to review its output against the document, and this sometimes catches errors. It is not a substitute for your check, because the same tendencies that produced the error may affect the review.

Can a second AI tool check the first? A second tool may catch different errors, and some teams find this useful. It does not remove the need for a person to check the output that will be relied on.

What if I find an error after the output has been used? Correct it as you would any error in your work: tell whoever relied on it, give the correct information and note what happened in your log. Then improve the prompt or the check so it does not recur.

Does checking get faster over time? Usually, yes. As you learn where a tool is reliable on a task, your checks become more focused. Keep the core steps, especially numbers, quotes and qualifications, even when the tool has been reliable for a while.

Checklist: checking an output

- Confirm every requested field or section is present.
- Check every number and date against the source.
- Confirm every quoted passage appears word for word.
- Check each important clause for its exceptions and conditions.
- Search the source for anything the output says is missing.
- Remove labels and conclusions the source does not support.
- Log the errors found and what caught them.

Contract Review with AI

For most small legal teams, contract review is where AI pays off first. The volume is high, the documents repeat, and the slowest part is often the first read. This chapter sets out a review workflow that uses AI for the first pass and keeps your judgment where it belongs.

What AI does well in review

AI tools are good at reading quickly and organizing what they read. They can extract the key terms of an agreement into a structured summary, find every clause on a given topic, compare an agreement against a list of positions, and spot clauses that look unusual compared with typical drafting. They are also useful for drafting: suggesting replacement wording, writing a reply to a counterparty's comment or explaining a clause in plain terms.

They are weaker at knowing what matters to your organization, at reading a clause in light of the rest of the deal and at knowing what the counterparty will accept. Those are the parts of review that depend on you.

A four-step review workflow

Step 1: Orient. Before using any tool, spend two minutes with the document. What is it, who is it with and which side are you on? What is the business trying to do? This context goes into your prompt and shapes your review.

Step 2: Extract. Ask the tool for a structured summary of the key terms, with section numbers and quoted wording (Chapter 4 gives a template). Check it using the method in Chapter 5. You now have a reliable map of the document.

Step 3: Compare. Give the tool your positions and ask it to list every departure. For each departure, ask for the clause wording, the departure in one sentence and a suggested response. Check each departure against the source, and look for departures the tool missed by scanning the clauses your positions cover.

Step 4: Decide and draft. Decide which departures matter, which you will accept and which you will push back on. Then use the tool to draft responses or replacement wording for the points you will raise. Review every draft before it goes anywhere.

Writing positions the tool can use

The comparison step depends on how clearly your positions are written. A position the tool can use has four parts: the topic, your standard position, your fallback and the trigger for escalation or walking away.

For an NDA, positions might cover the definition of confidential information, exclusions, term, permitted disclosures, return or destruction, remedies and any non-solicitation clause. For a services agreement, they might cover scope, acceptance, fees, liability cap, indemnities, intellectual property, data protection, termination and governing law.

Write positions in plain sentences. “Term of confidentiality: standard is 3 years from disclosure; fallback is 5 years; escalate if perpetual for all information.” The same positions guide your team, your outside counsel and the tool.

Common review errors to check

Beyond the general errors in Chapter 5, contract review produces some specific ones.

Carve-outs dropped from caps. The summary reports the cap and omits what sits outside it, such as confidentiality breaches, indemnities or data protection. This is one of the most common and most consequential errors.

Definitions not followed. A clause uses a defined term whose definition changes its meaning, and the summary reads the term in its ordinary sense.

Mutual and one-way confused. An obligation described as mutual applies to one party only, or the reverse.

Renewal and notice windows miscalculated. Auto-renewal clauses with notice windows are easy to summarize and easy to get wrong.

Order of precedence ignored. A schedule, order form or policy incorporated by reference overrides the main terms, and the summary reports the main terms as if they governed.

“Standard” without a basis. The tool describes a clause as standard or market. Remove this unless you have your own basis for it.

Drafting replies and redlines

AI can draft replies to counterparty comments, explain a position in plain terms and suggest replacement wording. Give it your position, the counterparty’s comment and the clause, and ask for a short reply. Review the reply for tone, accuracy and anything that concedes more than you intend.

When suggesting replacement wording, ask for the minimum change that achieves your position. Tools tend to rewrite more than necessary, and a heavily rewritten clause invites a heavily rewritten response. If your organization uses tracked changes, make the final edits in your document yourself, so the markup shows exactly what you changed.

What to send the business

Most business colleagues do not need your full analysis. They need to know whether they can sign, what they need to do and what risk they are carrying. Chapter 11 covers this in detail. For review, a short note with the decision, the two or three points that matter and any action required is usually enough.

When not to use AI for review

Some agreements are better reviewed without AI, at least for now. Agreements you are not permitted to share with the tool belong here. So do unusual, high-value agreements where the judgment is the whole job and the volume does not justify building a method. And so do documents in a format the tool handles poorly, such as scanned images with poor text recognition.

A worked example: an inbound SaaS agreement

The operations team at an invented company wants to sign a SaaS agreement for a scheduling tool. The vendor sends its standard terms and an order form. The in-house lawyer orients herself first: the company is the customer, the tool will hold employee names and schedules, the annual fee is modest and the business wants to start next month.

She runs the extraction step with her summary template. The summary covers the parties, fees, term, renewal, termination, liability, data and governing law, each with section numbers and quotes. Her check finds two problems: the summary reports a 12-month term without noting that the order form sets a 36-month term, which overrides the main terms under the order of precedence clause; and it reports the price as fixed when a section allows annual increases tied to the vendor's "then-current list price".

She then runs the comparison step against her SaaS positions. The tool lists six departures: the 36-month term; uncapped price increases; auto-renewal with a 90-day notice window; a liability cap at fees paid in the prior three months; no commitment to return data at termination; and a right for the vendor to use "aggregated and anonymized" customer data. Her scan of the clauses her positions cover finds a seventh the tool missed: the vendor may change its security practices on notice, without a commitment to maintain an equivalent standard.

She decides that three departures matter for this deal: the term, the price increases and the data return. She accepts the others with a note, given the modest fee and limited data. She asks the tool to draft minimum-change replacement wording for the three, reviews and edits the drafts, and makes the changes herself in the vendor's document with tracked changes.

Her note to the operations manager is four sentences: the agreement can be signed once three changes are agreed; the changes cap price increases, shorten the term to 12 months and require the vendor to return the company's data at the end; the markup is attached; please send it to the vendor and let her know if they push back.

A worked example: an NDA reply

A counterparty rejects a proposed change to an NDA's term, saying "a five-year term is market standard for an NDA of this kind." The lawyer's position is three years, with five as a fallback. She asks the tool for a short reply that accepts five years for trade secrets only and keeps three years for other information, and that answers the counterparty's point in its own terms.

The draft is clear but long, and it explains NDA terms generally. She cuts it to three sentences: the company accepts five years for trade secrets, which is where the longer protection matters; for other information, three years covers the period in which the information has value; and she proposes the split as a way to close the point. The tool saved her the first draft; the judgment and the cut were hers.

Questions people ask

Should I let the tool make tracked changes directly in the document? Some tools can. If you use this feature, review every change, and consider making final edits yourself so the markup reflects exactly what you intend. Tools sometimes make more changes than requested or change formatting in ways that obscure the substance.

What if I do not have written positions? Start with the agreement types you see most often. Write your standard, fallback and escalation point for the topics that come up in nearly every negotiation. A partial set of positions is far more useful than none, and it improves as you use it.

Can the tool tell me what the counterparty will accept? It can suggest what is common in drafting, but it has no reliable knowledge of what a particular counterparty will accept. That knowledge comes from your experience and the relationship.

Should I review the whole agreement myself as well? For agreements that matter, read the clauses your positions cover and any section the summary flags as unusual. For routine agreements, the structured summary and comparison, properly checked, may be enough under your team standard. Decide the level of review in advance, not under deadline pressure.

Checklist: reviewing an agreement with AI

- Note the document type, your side and the business goal before you start.
- Extract the key terms with section numbers and quotes, then check them.
- Compare against written positions and check each departure.
- Scan the clauses your positions cover for departures the tool missed.
- Check caps for carve-outs, definitions, renewal windows and order of precedence.
- Draft replies and wording with the tool; make final edits yourself.
- Send the business the decision, the key points and the action required.

Intake and Triage

Every small legal team knows the feeling of an inbox that is also a queue, a to-do list and a record of promises. Requests arrive by email, chat, hallway conversation and the occasional forwarded thread with "thoughts?" at the top. Before anyone can do the work, someone has to

work out what the request is, how urgent it is and what is missing. AI can help with that sorting, and it can help you build an intake process that reduces the sorting in the first place.

What triage involves

Triage answers four questions about each request. What kind of request is it (a contract review, a policy question, a dispute, an employment issue, a request for advice)? How urgent is it, and why? What information is missing before work can start? And who should handle it?

These questions are well suited to AI assistance because they follow patterns. Most requests fall into a handful of categories, urgency usually depends on a few signals (a signing date, a regulator, a customer escalation) and the missing information for each category is predictable.

Using AI to sort requests

A triage prompt gives the tool your categories, your urgency rules and the information each category needs. It then reads a request and returns a suggested category, urgency, missing information and a short draft reply asking for what is missing.

Write the categories and rules plainly. “Categories: commercial contract; NDA; employment; privacy or data; dispute or claim; corporate or governance; policy question; other. Urgent if: a signing or filing date within five business days; any regulator, court or law enforcement contact; any claim or threatened claim; any data incident. Missing information for a commercial contract: counterparty, value, our side, deadline, business owner, the document itself.”

The tool’s suggestions are a starting point. You decide the category and urgency, especially for anything that might be a claim or an incident, because those decisions have consequences the tool cannot weigh.

Confidentiality in triage

Intake material can be sensitive: complaints about colleagues, possible claims, personal data. Apply the sorting in Chapter 3 before using any tool. Where requests routinely contain sensitive material, consider using AI only on the structure of your intake process (the form, the categories, the replies) and not on the requests themselves until you have a tool cleared for that material.

Building a better intake form

The best triage is the triage you do not have to do. A short intake form, completed by the requester, can collect most of what you need before the request reaches you. AI is useful for drafting and testing that form.

Ask the tool to draft an intake form for your categories, with the minimum questions for each. Then test the form on a few recent requests: would the form have collected what you needed? Keep the form short. A form with twenty questions will be avoided; a form with six will be used.

A good intake form usually asks for: the requester and business owner; the type of request; the deadline and what drives it; the counterparty or people involved; the value or risk, in the

requester's words; and any documents. A free-text field for "anything else we should know" catches the rest.

Drafting replies

Many intake replies are the same: an acknowledgment, a request for missing information, a pointer to a template or policy, or a referral to the right person. AI can draft these quickly from your categories and rules. Keep a small library of approved replies and let the tool adapt them to each request.

Check each reply before sending. A reply that promises a turnaround you cannot meet, or answers a question you have not yet considered, creates a problem a slower reply would have avoided.

Tracking and reporting

Triage produces useful data: how many requests, of what type, from which teams, with what turnaround. Even a simple spreadsheet with one row per request gives you the material for the reports in Chapter 8 and the case for resources in Chapter 10. AI can help you summarize the tracker each month into a short note on volume, trends and bottlenecks.

Errors to check in triage

Urgency missed. A request that mentions a deadline in passing, or a regulator in a forwarded thread, is categorized as routine.

Category wrong for mixed requests. A request that is partly a contract question and partly an employment question is placed in one category, and the other part is lost.

Missing information invented. The draft reply assumes facts, such as a deadline or a counterparty, that the request did not give.

Tone mismatched. A reply to a stressed colleague reads as cold or bureaucratic.

A worked example: a Monday inbox

A legal assistant at an invented company starts the week with nineteen new requests. Using a triage prompt with the team's categories, urgency rules and missing-information lists, the assistant runs each request (with personal details removed) through the approved tool and gets a suggested category, urgency, missing information and a draft reply.

Fifteen suggestions look right on review. Four need changes. One request is labeled a routine NDA, but a line at the bottom of the forwarded thread mentions that the counterparty's lawyer has threatened to sue over a previous agreement; the assistant changes it to a possible claim and flags it to the lawyer immediately.

One request asks about both a contractor agreement and whether the contractor should be treated as an employee; the assistant splits it into two. One draft reply promises a review "by Wednesday", which the team cannot commit to; the assistant changes it to the team's standard

turnaround language. One request is in a language the tool translated loosely, and the assistant asks the requester to confirm the key points.

The whole sort takes a fraction of the usual Monday morning. More importantly, the possible claim reaches the lawyer within the hour instead of sitting in the queue.

Designing urgency rules

Urgency rules work best when they are specific and few. “Anything important” is not a rule. “A signing, filing or response date within five business days” is. Most small teams need five to eight rules. Common ones include: a deadline within a set number of days; any contact from a regulator, court or law enforcement; any claim, threatened claim or demand letter; any suspected data incident; any request from the chief executive or board; and any matter involving a named senior employee.

Review the rules after a month. If urgent requests keep slipping through, a rule is missing. If everything is urgent, the rules are too broad.

Questions people ask

Should the tool reply to requesters automatically? Not at first. Review every reply until you are confident the tool’s drafts are reliable for each category. Even then, keep a person in the loop for anything that might be a claim, an incident or a sensitive personnel matter.

What if requesters ignore the intake form? Make the form short, link it from where people already ask for help and reply to emails with a polite pointer to the form. Over time, show the benefit: requests through the form get faster answers because nothing is missing.

Can AI help us decide what not to do? It can help you apply rules you have written, such as a list of requests the business should handle with a template. The decision about what legal will and will not handle is a policy decision for the team and its leadership.

How do I handle requests that arrive in chat? Treat chat like email for triage purposes. Copy the request into your tracker, apply the same categories and rules, and reply with the same approved language. Keeping one record of requests matters more than the channel they arrive through.

Checklist: setting up triage

- Write your categories and urgency rules in plain sentences.
- List the information each category needs before work can start.
- Decide which requests are too sensitive for your current tools.
- Draft and test a short intake form against recent requests.
- Build a small library of approved replies.
- Track requests in one place, one row per request.
- Make the final call on category and urgency yourself.

Reporting to the Business

Legal teams report more than they think. There is the monthly update to the general counsel or chief executive, the board report, the status note on a dispute, the summary of a new law for a business team and the reply to “where are we on this?” that arrives every Thursday afternoon. Each takes time to assemble, and most follow a pattern. AI can draft them quickly from your notes, and it can help you write them in a way the reader will use.

What a good report does

A good report tells the reader what has changed, what they need to know and what they need to do. It puts the most important point first. It separates facts from judgments. And it is short enough that the reader will finish it.

These are habits, not features of a tool, and they are worth writing into your prompts. A tool asked to “write a monthly legal report” will produce something long and even; a tool asked for “the three changes that matter most this month, each in two sentences, with any decision needed at the top” will produce something a leader will read.

Building a report from your notes

Most reports start from material you already have: a tracker, meeting notes, emails and your own memory of the month. The work is sorting that material into what matters.

Give the tool your notes and a clear structure. A monthly report might use: decisions needed; what changed this month; what is coming next month; key numbers; risks to watch. Ask the tool to sort your notes into the structure, rank items by importance to the reader and list every figure with its source.

Then check the draft. Every figure should trace to a source in your notes. Every judgment (“this risk is increasing”) should be yours, not the tool’s. And anything the reader should not see, such as privileged analysis in a report that will circulate widely, should come out.

Writing for the reader

Different readers need different reports. A chief executive wants decisions and risks. A finance leader wants costs, savings and exposure in numbers. A board wants material risks, compliance status and anything requiring its approval. A business team wants to know what changes for them and what they need to do.

Tell the tool who the reader is and what they care about. Ask it to adjust the order, length and vocabulary. Then read the draft as that reader would, and cut anything they would skip.

Plain language without losing precision

Legal reports often fail by being either too technical or too vague. The aim is plain language that keeps the precision that matters. “We may have some exposure under the contract” is vague; “the contract caps our liability at fees paid in the last 12 months, about [AMOUNT], except for data breaches, which are uncapped” is plain and precise.

AI can help with this translation, but it tends to smooth out exactly the qualifications that matter. When you ask for plainer language, also ask the tool to keep every condition, exception and number. Check the result against the original analysis.

Status reports on matters

For disputes, investigations and other matters, a status report follows a stable pattern: what happened since the last report, what happens next and when, what decisions are needed, and any change in assessment or cost. AI can draft this quickly from your notes.

Take particular care with privilege and confidentiality. Status reports on disputes often contain legal advice. Decide who the report is for, mark it appropriately under your organization's practice and keep it out of tools not cleared for privileged material.

Summaries of new laws and guidance

Business teams often need short summaries of new laws, regulations or guidance. AI can draft these, but this is an area for extra care: tools may describe a law that has changed, a proposal that did not pass or a rule from a different jurisdiction.

Always work from the primary source. Give the tool the text of the law, regulation or guidance, and ask it to summarize only that text. Check the effective dates, the scope and any thresholds against the source. Add the date you checked and a note that the summary is general information.

Errors to check in reports

Figures without sources. A number appears that is not in your notes.

Judgments the tool made. The draft says a risk is “low” or a matter is “progressing well” without a basis in your notes.

Qualifications smoothed out. Plain language drops the exception that matters.

Stale or wrong law. A summary of a law reflects an older version, a proposal or another jurisdiction.

Too much. The draft includes everything in your notes, at even weight, and the important point is lost.

A worked example: the monthly report

The general counsel of an invented company sends a monthly report to the chief executive. It used to take most of a morning: pulling the tracker, reading notes, remembering what mattered and writing it up. Now she exports the tracker, pastes her rough notes and gives the tool a structure: decisions needed; three changes that matter most; coming up next month; numbers; risks to watch.

The first draft is well organized but treats everything as equally important, giving a routine NDA backlog the same weight as a regulatory inquiry. It also describes the inquiry as “unlikely

to result in significant liability”, which is a judgment she has not made. She asks the tool to rank items by importance to a chief executive and to remove any assessment not in her notes. The second draft leads with the regulatory inquiry and the decision she needs on outside counsel, and moves the backlog to the numbers section.

She checks every figure against the tracker, rewrites the risk section in her own words and removes a line about a dispute strategy that belongs in a privileged note rather than a report that will be forwarded. The finished report is one page. It takes her well under an hour, most of it spent on the parts that need her judgment.

A worked example: a new law for the sales team

A legal operations lead is asked to explain a new state privacy law to the sales team. He downloads the text of the law from the state legislature’s website and gives it to the tool with a request for a one-page summary of what changes for sales, limited to the text provided.

The draft is helpful but includes an effective date that does not match the text, apparently drawn from a different bill. Checking against the statute, he corrects the date, confirms the thresholds for which businesses are covered and adds a note that the summary is general information as of the date he checked. He also adds the one practical action the sales team needs to take, which the law implies but the draft did not state.

Questions people ask

Can the tool pull figures from our systems directly? Some tools can connect to spreadsheets and other systems. If you use this, check that the figures match the source and that the tool has access only to what it needs. Many teams find it simpler to export the figures themselves and paste them.

How do I avoid reports that sound machine-written? Edit the draft in your own voice, especially the opening and the risk section. Remove generic phrases and replace them with specifics. A report that uses your organization’s terms and refers to real decisions reads as yours.

Should I tell readers I used AI to draft the report? Follow your organization’s policy. Whatever the policy, the report is yours: you checked it and you stand behind it.

How often should reports go out? That depends on the reader and the volume of work. A short, regular report that leadership reads is more valuable than a long, occasional one. AI makes regular reporting easier to sustain.

Checklist: writing a report with AI

- Name the reader and what they care about.
- Give the tool a structure with the most important point first.
- Ask for every figure to be listed with its source.
- Keep every judgment your own.
- Ask plain language to keep every condition, exception and number.
- Work from primary sources for any summary of law or guidance.

- Read the draft as the reader would, and cut what they would skip.

Building a Team Standard

When one person uses AI well, the team gets faster in one place. When everyone uses it differently, the team gets faster in some places and riskier in others, and nobody can say exactly how the work is being done. A team standard fixes this. It does not need to be long, and it does not need a committee. It needs to say which tools, for which tasks, with which checks, and how the team learns.

Why a standard matters for a small team

Small teams often resist standards because they sound like bureaucracy. In practice, a standard saves time. It means each person does not have to work out the method from scratch, new staff can learn it in an afternoon and the team lead can answer the question “how does legal use AI?” with a document rather than a shrug.

A standard also protects people. If something goes wrong, a team member who followed the standard did what the team agreed. A team member who improvised is on their own.

What a team standard covers

A useful standard for a small legal team covers six things.

- 1. Approved tools and plans.** Which AI tools the team may use, on which plans, and for which kinds of information, drawing on the answers from Chapter 3.
- 2. Approved tasks.** Which tasks the team uses AI for, and which it does not. Start with the tasks you have tested and add others as they prove reliable.
- 3. Prompts and templates.** The tested prompt for each approved task, stored where everyone can find it, with the date it was last tested.
- 4. Checks.** The “what to check” list for each task (Chapter 5) and the level of review required before output is used.
- 5. Records.** What the team records: at a minimum, which tasks used AI, and the checking log for each approved task.
- 6. Learning.** How the team shares what it learns: a short monthly discussion, a shared note of errors found or a single owner who updates prompts and lists.

Writing it

Keep the standard short. Two or three pages is enough for most small teams, with the prompts and checklists in appendices or a shared folder. Write it in plain language, as instructions a new team member could follow on their first day.

AI can help draft the standard. Give it the six headings, your approved tools and tasks, and your existing checklists, and ask for a first draft. Then edit it so it describes what your team will

actually do, not what an ideal team might do. A standard nobody follows is worse than none, because it creates a record of a practice that does not exist.

Roles and supervision

In a mixed team, think about who does what. A paralegal or legal operations specialist may run the first pass and the structured checks; a lawyer may make the judgments and sign off. Write this down, so the standard makes clear where a lawyer decides.

Professional guidance on supervision applies to AI use. Lawyers who supervise others are generally responsible for making sure the work meets professional standards, including work done with AI tools. A clear standard is one of the most practical ways to meet that responsibility.

Fitting the company's policy

Many organizations now have an AI policy for all staff. The legal team standard should fit inside it, not contradict it. If the company policy is silent on something the legal team needs, such as privileged material, the team standard can be more specific. If the legal team is writing the company policy, the team standard is a useful test of whether the policy works in practice.

Introducing it

Introduce the standard in a short meeting. Walk through one task end to end: the prompt, the output, the checks and the record. Ask each person to try it on their next suitable task and bring one observation to the next meeting.

Expect the standard to change. The first version will have gaps, and tools will change. Set a date to review it, every three or six months, and update the prompts and checklists when you find errors or when a tool changes.

Keeping it current

AI tools change often, sometimes without notice. A prompt that worked last quarter may behave differently after a model update. Assign one person to retest the approved prompts when a tool announces a major change, and record the result. This is a small task that prevents a large surprise.

A worked example: a three-person standard

An invented legal team has a general counsel, a commercial lawyer and a paralegal. Each has been using the company's AI tool in their own way. The general counsel drafts board materials; the commercial lawyer reviews NDAs and vendor agreements; the paralegal summarizes contracts for the contract database. Nobody knows quite what the others do.

The general counsel writes a two-page standard. It lists the approved tool and plan, and the information that may be used with it. It lists four approved tasks: NDA review, vendor agreement summaries, contract database summaries and first drafts of board materials. For each, it names the prompt file, the "what to check" list and who signs off: the commercial

lawyer for NDAs and vendor summaries, the general counsel for board materials, and the paralegal for database summaries with a monthly spot check by the commercial lawyer.

It sets out three rules that apply to everything: no privileged communications in the tool; personal data removed unless a task needs it; and every output checked against the source before use. It names the commercial lawyer as owner of the prompts, responsible for retesting after major tool changes. And it sets a quarterly 30-minute meeting to share errors found and update the lists.

At the first meeting, the paralegal reports that the database summaries regularly misread renewal dates in agreements where the order form overrides the main terms. The commercial lawyer adds an order-of-precedence check to the prompt and the list. One error, found by the person closest to the work, improves the method for everyone.

What to leave out

A small team's standard should leave out anything the team will not do. It should not promise reviews nobody has time for, records nobody will keep or meetings nobody will attend. It should not restate the company's AI policy at length; a reference is enough. And it should not try to cover every possible task. A standard that covers four tasks well is more useful than one that covers twenty in general terms.

Questions people ask

Who should own the standard? The team lead usually owns the standard itself. Day-to-day upkeep of prompts and lists can sit with whoever uses them most, often a lawyer or legal operations specialist.

How do we handle a team member who prefers not to use AI? A standard can approve tasks without requiring everyone to use AI for them. What matters is that anyone who does use AI for an approved task follows the standard.

What if a team member finds a better prompt? Test it on the team's test documents, and if it performs better, update the shared prompt and note the change. The standard should make improvement easy.

Should the standard cover personal use of AI for learning? It can include a short note that learning and experimentation on invented or public material is encouraged, which helps people build skills without risk.

Checklist: building a team standard

- List approved tools, plans and the information each may be used with.
- List approved tasks, starting with those you have tested.
- Store each tested prompt with its test date and "what to check" list.
- Set the level of review required for each task, and who signs off.
- Decide what the team records and where.
- Introduce the standard with one worked task, end to end.
- Set a review date and an owner for retesting after tool changes.

Measuring Time Saved Honestly

Time is why most people try AI, and it is what most budget holders ask about. “How much time does this save?” is a fair question, and a vague answer weakens your case. An inflated answer is worse, because it sets an expectation you will later have to walk back. This chapter sets out a simple, honest way to measure time saved on your own work.

Why general claims are not enough

You will see many claims about how much time AI saves lawyers. Some come from surveys of what people expect or believe, some from vendor case studies and some from controlled studies of particular tasks. Each may be accurate for what it measured. None measures your work, with your documents, your tools and your standards.

Your own measurement is more useful than any general claim, for two reasons. It reflects the work you actually do, including the review time that general claims often leave out. And it is evidence you can explain: you know exactly how it was measured and what it covers.

What to measure

Measure the whole task, from start to a result you would rely on. For a contract summary, that means the time from opening the document to a summary you would send, including writing or adapting the prompt, running it, checking the output and fixing errors. Leaving out the checking time makes the saving look larger and the measurement less honest.

Measure the same task both ways: once without AI, once with it. Use comparable documents, so the comparison is fair. If you can, have the same person do both, on different but similar documents, so familiarity with one document does not distort the other.

A simple timing method

Pick one task and a set of comparable documents, ideally invented or public ones you can reuse. Five to ten documents is enough to see a pattern.

For each document, record:

- the date and the tool and model used;
- who did the work;
- time without AI, from start to a result you would rely on;
- time with AI, from start to the same standard, including checking;
- errors found in the AI output and how long they took to fix.

Use a simple timer and round to the nearest minute. Then take the median of each column. The median is less affected by one unusually long or short run than the average.

Reporting what you found

State the result with its basis. “In our tests on eight sample services agreements in March, a first-pass summary took a median of [X] minutes with AI, including checking, against [Y]

minutes without. Your time will depend on your documents.” This sentence is specific, honest and hard to argue with.

Avoid turning a small test into a large claim. Eight documents show a pattern for that task; they do not show that AI saves a certain percentage of legal time generally. If you want to estimate annual savings, multiply the per-task saving by the number of such tasks you actually handle, and say that you did so.

When the saving is small or negative

Sometimes the measurement shows little or no saving, or shows that AI is slower once checking is included. That is a useful result. It may mean the prompt needs work, the task is not well suited to AI or the checking can be made more efficient. It may also mean the task is better left as it is.

Report these results too. A team that reports only successes is less credible than one that says “we tested three tasks; two saved time and one did not, so we dropped it.”

Time is not the only measure

Time saved is the easiest benefit to measure, but it is not the only one. AI can make work more consistent, because the same prompt and checklist produce the same structure each time. It can reduce the backlog, because first passes happen sooner. And it can shift time from routine work to judgment work, which is harder to count but often more valuable.

Note these benefits where you see them, but be careful about claiming them without evidence. “Summaries now follow the same structure every time” is a fact you can show. “Quality has improved” is a judgment that needs support.

Using the results

Your measurements are useful in three places. In your team standard (Chapter 9), they show which tasks are worth the effort. In budget requests, they give finance a specific, explained figure in place of a general claim. And in your own planning, they show where your time goes and where it could go instead.

Keep the records. If anyone later asks how you reached a figure, you can show them.

A worked example: timing NDA review

The lawyer at the invented Harbor Lane Software times NDA review on eight invented NDAs, four reviewed each way, then the reverse on a second set of eight so each NDA is reviewed both ways by different sessions. She records the time from opening the document to a completed issues list and reply, including writing the prompt for the first one, running it, checking the output and fixing errors.

Her records show that the time with AI varies more than the time without it: some NDAs need only a quick check, while the badly drafted one takes longer than doing it by hand because the output needs heavy correction. The median is what she reports, with the number of runs and the date. She also records that the time without AI includes several minutes of “getting started”

on each document that the AI-assisted version largely avoids, which matches her sense of where the saving comes from.

She writes the result as a single sentence with its basis, in the form this chapter recommends, and files the timing sheet with her test documents. When she later asks for budget for a team plan, the sentence goes into her request, and the timing sheet is ready if finance asks how she measured it.

Common measurement mistakes

Leaving out the check. Timing only the AI's run, not the review, makes the saving look larger than it is.

Comparing different documents. Timing a simple agreement with AI and a complex one without it says nothing about AI.

Counting the first try. The first run of a new prompt includes learning time. Note it, but do not treat it as typical.

Extrapolating too far. A saving on one task does not imply the same saving on other tasks.

Forgetting the date and model. Tools change. A result without a date and model is hard to rely on later.

Questions people ask

Can I use published figures instead of measuring? Published figures can give context, quoted with their source and described as that source's finding. They do not measure your work, and a request supported by your own measurement is stronger.

What if I do not have time to run a careful test? Even three or four timed runs, recorded honestly, are better than none. State the small number when you report the result.

Should someone else time my work? If you can arrange it, an independent tester strengthens the result. For most small teams, careful self-timing with honest records is a reasonable start.

How do I measure time saved for the whole team? Measure task by task, then multiply each task's saving by how often the team does that task, and say that is how you calculated it. Avoid a single percentage for all legal work.

Checklist: measuring time saved

- Pick one task and five to ten comparable documents.
- Time the whole task both ways, including prompt writing and checking.
- Record date, tool, model, person, times and errors.
- Take the median of each column.
- State the result with its basis and a note that results vary.
- Report tasks where AI did not save time, as well as those where it did.
- Keep the records for as long as you use the figure.

Explaining Legal to the Business

Much of a small legal team's value depends on whether the business listens. Advice that is right but not read, or read but not understood, does little. AI will not fix a relationship, but it can help with the part that depends on wording: turning analysis into something a busy colleague will read, understand and act on. This chapter covers how.

Why legal advice gets lost

Legal advice tends to get lost for predictable reasons. It arrives too late, after the decision is effectively made. It is too long, and the point is in the fourth paragraph.

It uses terms the reader does not know. It explains the law without saying what to do. Or it reads as a reason to say no, when the reader was looking for a way to say yes.

Most of these are problems of structure and wording, and those are things AI can help with.

The shape of a useful note

A useful note to a business colleague has a simple shape. The first sentence gives the answer or decision. The next two or three sentences give the reasons that matter to the reader. Then comes what the reader needs to do, and by when. Detail, if any, follows for those who want it.

"You can sign this agreement once two changes are made. The liability cap is lower than our standard, and the renewal clause would lock us in for another year unless we give notice by March. I have drafted the changes; please send them to the vendor this week. More detail is below if useful."

That note takes thirty seconds to read and leaves no doubt about what happens next.

Using AI to shape the note

Give the tool your analysis and ask it to produce a note in that shape. Tell it who the reader is, what they care about and what they need to do. Ask it to keep every condition, number and date, and to avoid legal terms unless they are explained.

Then check the draft against your analysis. Plain language has a way of dropping the qualification that matters. And check the tone: a note that is clear but curt can damage the relationship as much as one that is unclear.

Explaining caution without blocking

Business colleagues sometimes read legal caution as obstruction. One way to change this is to explain the reason for the caution in the reader's terms. "This clause would let the vendor raise prices by any amount each year" is more persuasive than "this clause is unfavorable."

Another is to offer a way forward with every concern. "We can accept this if the increase is capped at five percent" turns a no into a negotiation. AI can help you draft these alternatives, but the judgment about what is acceptable is yours.

Explaining AI use itself

You may also need to explain how legal uses AI. Some colleagues will worry that legal is cutting corners; others will want legal to use AI more. A short explanation helps both groups: which tasks legal uses AI for, how the output is checked and what stays with a lawyer. Your team standard (Chapter 9) is the source for this explanation.

Templates for common notes

Many notes repeat. Keep a few templates:

- **Sign or do not sign:** decision, reasons, changes needed, action and deadline.
- **What this new rule means for you:** what changes, who is affected, what to do and by when, where to ask questions.
- **Status update:** what happened, what happens next, any decision needed.
- **Why we need more time:** what is outstanding, why it matters, the revised date.

AI can adapt these to each situation, and the template keeps the structure consistent.

Errors to check

Qualifications dropped. The plain version loses the exception, condition or number that matters.

Overstated certainty. “You can sign” when the honest answer is “you can sign if the vendor accepts our changes.”

Advice the analysis does not support. The draft recommends something you did not conclude.

Tone. The draft is too casual for the reader, or too formal for the relationship.

A worked example: from memo to message

A commercial lawyer at an invented company has written a two-page internal memo on a distribution agreement. The memo is thorough: it covers exclusivity, minimum purchase commitments, termination, intellectual property and governing law. The sales director who asked for the review needs to know one thing: can she agree to the exclusivity the distributor wants?

The lawyer gives the tool the memo and asks for a note to the sales director, answering her question in the first sentence, with only the points that bear on exclusivity and the action she needs to take. The first draft answers the question but adds a paragraph on governing law that the sales director does not need. It also says “exclusivity is acceptable”, dropping the condition in the memo that exclusivity should end if the distributor misses its minimum purchases.

The lawyer cuts the governing law paragraph and restores the condition. The final note reads: “You can agree to exclusivity for the western region, provided it ends if the distributor misses its minimum purchases in any two consecutive quarters. Without that condition, we would be

locked into one distributor even if sales fall. I have drafted the wording; please propose it in your next call. The full memo is on file if you want the detail.”

A worked example: a difficult no

A business manager wants to sign a vendor’s terms today. The terms include an uncapped indemnity from the company to the vendor for any claim arising from the company’s use of the service. The lawyer’s answer is no, at least not as drafted.

She asks the tool for a short note explaining the problem in business terms and offering a way forward. She edits the draft so it reads: “Please hold off signing today. As written, we would pay the vendor’s costs for any claim connected with our use of the service, with no limit, including claims that are the vendor’s own fault. If the vendor will limit this to claims caused by our misuse, and cap it at the contract value, we can sign. I have sent the wording to the vendor’s contact and copied you.”

The note says no without sounding obstructive, because it explains the risk in the manager’s terms and shows the path to yes.

Questions people ask

How do I keep my own voice when using AI to draft notes? Edit the opening and the closing yourself, and use the words you would normally use with that colleague. Over time, you can give the tool a few examples of your notes so its drafts start closer to your style.

What if the business wants everything in one line? Some decisions fit in one line, and that is fine. Where a condition matters, keep it, even if the note becomes two lines.

Should I send the business the full analysis? Usually not, unless they ask. Keep the analysis on file and offer it. Most readers want the decision and the action.

How do I handle disagreement with the business? State the risk plainly, offer alternatives and say who decides. AI can help you draft a balanced note, but the relationship and the decision belong to the people involved.

Checklist: writing a note for the business

- Put the answer or decision in the first sentence.
- Give only the reasons that matter to this reader.
- Say what the reader needs to do and by when.
- Explain caution in the reader’s terms, with a way forward.
- Check that every condition, number and date survived.
- Read the draft for tone, as the reader would.
- Keep templates for the notes you write most often.

What Comes Next: Agents

Everything in this book so far uses AI as an assistant: you give it a task, it gives you a draft, you check it. The next step, already arriving in many tools, is the agent: an AI system that carries out a sequence of steps on its own, such as reading a contract, comparing it to a playbook, drafting a reply and filing the result. Agents promise more time saved, and they bring new questions about control. This chapter explains what agents are, where they fit for a small legal team and how to use them without losing control of the work.

What an agent is

An agent is an AI system that can plan and carry out several steps toward a goal, often using other tools along the way: reading files, searching, writing documents, sending messages. Where an assistant answers one prompt, an agent works through a task, deciding the next step from the result of the last one.

For legal work, an agent might take an incoming NDA from an inbox, check it against your playbook, draft a markup and a reply, and place both in a folder for your review. Each step is something this book has already covered; the difference is that the agent links them together.

Where agents fit

Agents fit best where a task is well defined, repeats often and has a clear point where a person reviews the result. The contract review workflow in Chapter 6 is a good example: an agent can run the extraction and comparison steps, and you make the decisions and approve the drafts.

Agents fit less well where the task depends on judgment at many points, where errors are hard to see or where the agent would need access to systems or information it should not have.

Control points

The most important design question for any legal agent is where a person approves the work. A well-designed agent stops at defined points and waits for approval before doing anything with consequences: sending a message, signing, filing, or sharing a document outside the team.

Write these control points down before using an agent. “The agent may read incoming NDAs and draft a markup and reply. It may not send anything. A lawyer reviews every draft before it is sent.” This is the agent version of the team standard in Chapter 9.

Access and permissions

An agent can only do what it has access to, which makes access the main safeguard. Give an agent the least access it needs: read access to the folder where contracts arrive, write access to a drafts folder, no access to send email or to systems it does not need. Review the access regularly, and remove it when the agent is no longer used.

The confidentiality questions in Chapter 3 apply with more force to agents, because an agent may handle more information, more often, with less direct attention. Know which tool runs the agent, on which plan, under which terms, before connecting it to anything sensitive.

Testing an agent

Test an agent as you would test a prompt, but over the whole sequence. Run it on a set of invented documents that includes easy cases, hard cases and cases designed to trip it up: a missing clause, a contradictory definition, an unusual governing law. Check every step, not just the final output, because an error early in the sequence can carry through to a result that looks reasonable.

Record what the agent gets right, what it gets wrong and where it stops for approval. Retest when the underlying model or tool changes.

Errors specific to agents

Errors that compound. A misread at step one shapes every later step, and the final output looks consistent because it is consistently wrong.

Actions without approval. An agent configured too loosely takes a step, such as sending a message, that should have waited for a person.

Scope drift. An agent asked to do one task does more, such as editing a document it was asked only to read.

Silent failure. An agent fails partway through and produces an incomplete result without saying so.

A sensible path for a small team

You do not need to build agents to benefit from them. Many general AI tools now include agent features, and legal technology providers offer agents for common tasks. The sensible path is the same one this book has followed throughout: start with one well-defined task, test it on safe material, write down the control points and the checks, and expand only when the method works.

If you have built good prompts, clear playbooks and reliable checklists for the tasks in this book, you already have most of what a legal agent needs. The agent links the steps; your standards make the result something you can stand behind.

A worked example: an NDA agent with one control point

The invented Harbor Lane Software adopts an agent for NDA review once the lawyer's manual method is reliable. The agent watches a shared folder where the sales team saves incoming NDAs. For each new file, it extracts the key terms, compares them against the lawyer's written positions, drafts a markup and a reply, and saves all three to a review folder with a short note of the departures it found.

The control point is simple: the agent cannot send anything. The lawyer reviews each draft, edits as needed and sends it herself. The agent has read access to the incoming folder and write access to the review folder, and nothing else.

Before going live, the lawyer tests the agent on her twelve invented NDAs, including the hard cases from her original testing. It handles most well but produces an incomplete markup on one NDA saved as a scanned image, without saying the text could not be read. She adds a rule that the agent must flag any document where it cannot read the full text, and she retests. She also records the agent's results, model and date, and sets a reminder to retest after the next major update to the tool.

The result is that NDAs reach her already summarized and marked up, and her time goes to the departures that matter. The judgment stays with her, and so does the decision to send.

Questions to ask a vendor about a legal agent

If you are considering an agent from a vendor, ask: what steps the agent takes on its own; where it stops for approval and whether that can be changed; what access it needs and whether that can be limited; where data goes and under which terms; how it reports failures; how it has been tested on legal tasks and whether you can see the results; and what happens to your data and configuration if you stop using it.

The answers will tell you much about whether the agent fits a careful legal team.

Questions people ask

Will agents replace legal staff? Agents change the work, moving time from routine steps to review and judgment. For small teams, the more realistic effect is that the same people handle more work with more consistency, and spend more of their time on the work that needs them.

Is it safe to let an agent send routine replies? Only after extensive testing, for a narrow category, with clear rules and regular spot checks. Many teams will reasonably decide to keep a person approving every outgoing message.

How do agents fit with professional responsibility? The same duties apply: competence, confidentiality, supervision and communication. A well-defined agent with control points and records is easier to supervise than ad hoc use of AI tools.

Checklist: before using an agent

- Choose one well-defined, repeating task with a clear review point.
- Write down what the agent may and may not do, and where a person approves.
- Give the agent the least access it needs, and review it regularly.
- Confirm the tool, plan and terms before connecting it to sensitive material.
- Test the whole sequence on invented documents, including hard cases.
- Check every step, not just the final output.
- Retest when the underlying model or tool changes.

A closing word

Using AI in legal work well is not a matter of finding the perfect tool. It is a matter of applying the habits good legal teams already have: clear instructions, careful review, written standards

and honest measurement. Those habits make AI useful, and they make the result something you can stand behind. Start with one task. The rest follows.

Appendix A: A Starter Prompt Library

These prompts follow the five-part structure in Chapter 4. Each is a starting point: adapt the bracketed parts, test it on invented documents in the tool you will use and add your own “what to check” list before relying on it. None of these prompts has been tested by ProvisoIQ for this book; treat each as untested until you have run it yourself.

A1. Agreement summary for a business reader

“You are assisting an in-house lawyer at [TYPE OF ORGANIZATION]. Prepare a summary of the [AGREEMENT TYPE] between the markers below for [READER], who needs to decide [DECISION]. Using only the agreement, complete these fields: parties and roles; what is being provided; term, renewal and notice; fees and how they can change; liability cap, what is excluded from it and any exceptions; indemnities; termination rights; data and confidentiality; governing law and disputes; anything unusual. For each field, give the section number and quote the key words. Write ‘Not addressed’ where the agreement is silent. Do not describe any term as standard, reasonable or favorable. If you are unsure how to read a clause, say so and quote it. Keep the summary under 400 words.”

What to check: every number and date; cap exceptions; renewal notice windows; order of precedence between the main terms and any order form or schedule.

A2. Issues list against your positions

“You are assisting an in-house lawyer reviewing an inbound [AGREEMENT TYPE]. Our side is [OUR SIDE]. Our positions are listed below, each with a standard, a fallback and an escalation point. Compare the agreement between the markers to these positions. Return a table with columns: section, topic, clause wording (quoted), departure from our standard (one sentence), whether it falls within our fallback (yes, no or unclear) and suggested response. Include a row for any position the agreement does not address. Do not assess legal risk beyond the positions given. If you are unsure whether something is a departure, include it and explain why.”

What to check: departures missed in clauses your positions cover; definitions that change a clause’s meaning; mutual and one-way obligations.

A3. Reply to a counterparty comment

“We are negotiating [AGREEMENT TYPE] on behalf of [OUR SIDE]. The counterparty has commented on [CLAUSE] as follows: [COMMENT]. Our position is [POSITION], and we can offer [CONCESSION, IF ANY]. Draft a reply of no more than four sentences that answers the counterparty’s point in its own terms, states our position and, if a concession is available, offers it as the way to close the point. Do not restate the clause or explain the law. Use a courteous, direct tone.”

What to check: anything conceded beyond what you intend; tone; whether the reply answers the question actually asked.

A4. Minimum-change wording

“Here is a clause from an agreement: [CLAUSE]. Our position is [POSITION]. Suggest the smallest change to the clause’s wording that achieves our position. Show the original and the revised wording, and list each change. Do not rewrite parts of the clause that do not need to change.”

What to check: unrequested changes; defined terms used correctly; cross-references still accurate.

A5. Intake triage

“You are helping a legal team sort incoming requests. Categories: [CATEGORIES]. Urgent if: [URGENCY RULES]. Information needed by category: [LIST]. For the request below, return: suggested category (or categories, if the request covers more than one); urgent or not, with the reason; information missing; and a short, polite reply asking for the missing information without promising a turnaround. Do not assume any fact the request does not state.”

What to check: urgency signals buried in forwarded threads; mixed requests; replies that promise dates.

A6. Monthly report from notes

“You are helping a general counsel write a monthly report for [READER]. Using only the notes and figures below, produce a report with these sections: decisions needed; the three changes that matter most this month; coming up next month; key figures; risks to watch. Rank items by importance to [READER]. List every figure with its source in the notes. Do not add any assessment, prediction or characterization not stated in the notes. Keep it to one page.”

What to check: figures without sources; assessments the tool added; privileged material in a report that will circulate.

A7. Plain-language note for the business

“Here is my analysis: [ANALYSIS]. Write a note to [READER], who asked [QUESTION]. Put the answer in the first sentence. Give only the reasons that matter to this reader, in plain terms. Say what they need to do and by when. Keep every condition, exception, number and date from the analysis. Keep it under [LENGTH] words.”

What to check: dropped conditions; overstated certainty; tone for the relationship.

A8. Summary of a new law from its text

“Below is the text of [LAW OR GUIDANCE] from [SOURCE], retrieved on [DATE]. Using only this text, summarize what it requires of [TYPE OF ORGANIZATION OR TEAM], who is covered and any thresholds, the effective dates and any exceptions. Quote the provisions you rely on with

their section numbers. Do not describe any other law, proposal or guidance. If the text is unclear, say so.”

What to check: effective dates; thresholds; whether the text you supplied is the current, enacted version.

A9. Policy first draft

“Draft a first version of a [POLICY TYPE] for [TYPE OF ORGANIZATION] with these features: [FEATURES]. Use plain language and short sections. Include headings for purpose, scope, rules, responsibilities, exceptions and review. Mark any point that depends on a legal requirement as [CHECK LAW] rather than stating the requirement. Do not cite laws.”

What to check: anything stated as a legal requirement; scope that is broader than intended; rules nobody can follow.

A10. Checking prompt

“Below are a document and a summary of it. Check the summary against the document. For each statement in the summary, say whether the document supports it, quoting the supporting wording, or whether it is unsupported or inaccurate. List anything important in the document that the summary omits, especially exceptions, conditions, numbers and dates.”

What to check: use this as an extra pass, not a replacement for your own check.

A11. Test document generator

“Create an invented [AGREEMENT TYPE] of about [LENGTH] words between two invented companies, for testing a review method. Include these features: [FEATURES, SUCH AS A CAP WITH CARVE-OUTS, AN AUTO-RENEWAL WITH A NOTICE WINDOW, AN INCONSISTENT DEFINED TERM]. Do not use real company names. After the agreement, list each feature and the section where it appears.”

What to check: that the features are present where the list says; that no real names appear.

A12. Prompt improvement

“Here is a prompt I use for [TASK]: [PROMPT]. Here are errors it produced in testing: [ERRORS]. Suggest changes to the prompt that would reduce these errors, explaining each change in one sentence. Do not change parts of the prompt unrelated to the errors.”

What to check: retest any revised prompt on your full test set before adopting it.

Appendix B: A Sample Team Standard

This sample shows the shape of a short team standard for an invented three-person legal team. Replace every bracketed item, and keep only what your team will actually do.

[Team name] AI Use Standard

Version [NUMBER], [DATE]. Owner: [NAME OR ROLE]. Next review: [DATE].

1. Purpose

This standard sets out how the legal team uses AI tools, so that our work is consistent, our checks are reliable and anyone can explain how we use AI. It sits within the company's AI policy, [POLICY NAME AND LOCATION], and where the two differ, the company policy governs unless this standard is stricter.

2. Approved tools

We use [TOOL] on the [PLAN] plan, under the company's agreement dated [DATE]. Training on our input is [OFF OR ON], confirmed on [DATE]. This tool may be used with: public and invented material; company confidential material, including draft agreements with counterparty names. It may not be used with: privileged communications or legal advice; personal data beyond what a task requires; material whose confidentiality terms forbid disclosure to service providers. No other AI tool is used for work material without the owner's approval.

3. Approved tasks

Task	Prompt file	What to check list	Who signs off
NDA review	[FILE]	[FILE]	[ROLE]
Vendor agreement summary	[FILE]	[FILE]	[ROLE]
Contract database summary	[FILE]	[FILE]	[ROLE], with monthly spot check by [ROLE]
Board materials first draft	[FILE]	[FILE]	[ROLE]

Any other task is used only after testing on invented documents and the owner's approval, recorded in the task table.

4. Rules for every task

Every output is checked against its source before use, following the task's "what to check" list. Section numbers and quoted wording are required for any claim about a document. Personal data is removed before input unless the task requires it. Outputs are never sent outside the team without review by the person who signs off for that task.

5. Records

For each approved task, the person who signs off keeps a checking log: date, tool and model, document reference, errors found and what caught them. Timing records for any time-saved figure are kept with the test documents for as long as the figure is used.

6. Upkeep and learning

The owner retests approved prompts on the test documents when the tool announces a major change, and records the result. The team meets for 30 minutes each [MONTH OR QUARTER] to share errors found and agree updates to prompts and lists. Anyone may propose a change to a prompt; changes are adopted after testing on the test documents.

7. Questions

Questions about this standard go to [OWNER]. Questions about the company AI policy go to [POLICY OWNER].

Appendix C: Glossary

These are working definitions for legal teams, written in plain terms. Providers sometimes use the same words differently, so check how a term is used in the terms that apply to you.

Agent. An AI system that carries out several steps toward a goal, deciding each next step from the last result and often using other tools along the way.

Approved tool. An AI tool your organization has cleared for work use, on a named plan, for named kinds of information.

Business plan. A paid plan for organizations, usually with different data terms from a consumer plan. Names vary by provider: team, business, enterprise.

Consumer plan. A free or individual paid plan intended for personal use. Its data terms may permit uses of input that a business plan does not.

Context window. The amount of text a model can consider at once. Long documents may exceed it or strain it, which can cause later parts of a document to be handled less reliably.

Control point. A step in a workflow where a person must review or approve before anything continues, such as before a message is sent.

Data processing addendum. A contract term, often attached to a provider's agreement, that sets out how the provider handles personal data on the customer's behalf.

Fine-tuning. Training a model further on a particular set of data. Not something most small teams need to do.

Generative AI. AI that produces new text, images or other content in response to instructions.

Hallucination. A common term for AI output that states something false or unsupported as if it were true, such as a clause that is not in a document or a citation that does not exist. This book calls it invented content.

Input. What you give the tool: your prompt, any document and any follow-up questions.

Large language model. The type of model behind most general AI text tools, trained on large amounts of text to predict and produce language.

Model. The underlying AI system a tool uses. A single product may offer several models, and providers update models over time.

Output. What the tool gives back.

Playbook. A written set of positions for a type of agreement, usually with a standard, a fallback and an escalation point for each topic.

Prompt. The instruction you give an AI tool.

Prompt template. A tested prompt saved for reuse, with placeholders for the parts that change.

Retention. How long a provider keeps your input and output, and whether you can delete it or shorten the period.

Test set. A group of documents, usually invented, used to test a prompt or agent before relying on it, including easy cases, hard cases and cases designed to expose errors.

Training on input. A provider's use of what you enter to improve its models. Whether this happens depends on the plan, settings and terms.

What to check list. A short list of the errors a particular task tends to produce, used every time the output is reviewed.

Appendix D: A Thirty-Day Plan

This plan turns the book into four weeks of small steps for one person, at about two to three hours a week. If you lead a team, run it yourself first, then use your notes to guide the team through it.

Week 1: Prepare

Day 1. List ten tasks from your last two weeks and score them on the four questions in Chapter 2. Choose your first task.

Day 2. Find out which AI tools your organization has approved, on which plan, and for what kind of information. If you cannot find out, write the five questions from Chapter 3 and send them to the person who manages the tools.

Day 3. Write down what a good result looks like for your first task, in three or four sentences. If the task involves agreements, write your standard positions for the five topics that come up most.

Day 4. Create five invented test documents for the task, deliberately varied: one clean, one with a common problem, one with an unusual clause, one badly drafted and one long. Prompt A11 in Appendix A can help.

Day 5. Review the five test documents yourself, the way you normally would, and time each one. Keep your results; they are the benchmark.

Week 2: Test

Day 6. Write your first prompt using the five-part structure in Chapter 4. Run it on the first test document and check the output using Chapter 5.

Day 7. Run the prompt on the other four test documents. Record each error, what caught it and how long the whole task took, including checking.

Day 8. Improve the prompt to address the errors you found, changing one thing at a time. Rerun on all five documents.

Day 9. Write the “what to check” list for the task: five to eight items, drawn from the errors you found.

Day 10. Decide: adopt, adjust or move on. If the task is not working after three rounds, return to your list and choose the next candidate.

Week 3: Use

Day 11. If your tool is approved for real material of this kind, use the method on a real task, following the confidentiality steps in Chapter 3. If not, keep testing on invented material.

Days 12 to 14. Use the method on each suitable task that arrives. Keep the checking log for each one.

Day 15. Review the log. Are new errors appearing? Is checking getting faster? Update the prompt and the list.

Week 4: Record and share

Day 16. Write a one-page note of your method: the task, the tool and plan, the prompt, the “what to check” list and who signs off. This is the first entry in a team standard.

Day 17. Time the task on five more comparable documents, both ways, following Chapter 10. Write the result as one sentence with its basis.

Day 18. Tell your manager, in a short note, how you are using AI on this task, the checks you apply and the result of your timing.

Day 19. Choose your second task, using what you learned from the first.

Day 20. Set a reminder to retest your prompt when your tool announces a major change, and to review your method in three months.

After thirty days

By the end of the month, you will have one task you can do with AI reliably, a prompt and checklist you trust, honest timings and a written note of your method. That is a small, solid result, and it is the foundation for everything else in this book. The second task usually goes faster than the first, because the habits are already in place.

Appendix E: Building a Test Set

A test set is a small group of documents you use to test a prompt, a checklist or an agent before relying on it. A good test set tells you how a method fails, not only whether it works. This appendix explains how to build one and gives a short invented example.

Why invented documents

Invented documents have three advantages. They are safe: you can use them in any tool without confidentiality concerns. They are reusable: you can rerun the same set after a tool changes and compare results. And they are designed: you can include exactly the problems you want to test for, in known places, so you know whether the method caught them.

Public documents, such as published standard terms, can supplement an invented set, but check that you are permitted to use them this way, and remember that a tool may have seen common public documents before.

What to include

For each task, aim for five to twelve documents that cover:

- **A clean case.** A well-drafted document close to your standard, to check the method does not invent problems.
- **Common problems.** The departures you see most often, placed where they usually appear.
- **Hidden problems.** Issues placed where they are easy to miss: in a definition, a schedule, an order form or a cross-reference.
- **Structural traps.** An order of precedence clause that changes which terms govern; a defined term used inconsistently; a mutual clause that is one-way in one sentence.
- **Absent clauses.** A document missing a clause you would expect, to check the method reports the gap and does not invent the clause.
- **A long document.** To check the method holds up past the first few pages.
- **A badly drafted document.** With inconsistent numbering and loose wording, closer to some of what arrives in practice.

For each document, keep an answer key: the problems you placed and where. The key is what makes the test set useful, because it lets you score each run.

Scoring a run

For each document, record: the problems the method caught; the problems it missed; anything it reported that was not a problem; and the time taken, including checking. Across the set, these records show the method's reliability and the shape of its errors.

Update your “what to check” list from the misses, and your prompt from the false reports.

A short invented example

The extract below is from an invented NDA designed to test four features: a residuals clause, a term that differs between the main text and a definition, a one-way obligation described as mutual and a missing return-or-destruction clause.

MUTUAL NONDISCLOSURE AGREEMENT between Example Harbor Inc. (“Company”) and Sample Ridge LLC (“Recipient”).

1. Definitions. “Confidential Information” means all non-public information disclosed by Company to Recipient. “Term” means three years from the Effective Date.

2. Obligations. Each party shall hold the other’s Confidential Information in confidence and use it only to evaluate the proposed transaction.

3. Residuals. Nothing in this Agreement restricts Recipient’s use of information retained in the unaided memory of its personnel.

4. Duration. The obligations in Section 2 continue for five years after the Effective Date.

5. Governing law. This Agreement is governed by the laws of the State of [STATE].

Answer key: Section 1 defines Confidential Information as Company’s information only, so despite the title and Section 2, protection is one-way; Section 3 is a residuals clause that can largely undo the protection; Sections 1 and 4 give different durations, three years and five years; there is no return or destruction clause.

A method that reports this as a standard mutual NDA with a three-year term has failed on all four points. A method that catches all four, quoting the relevant sections, is working.

Keeping the set current

Add a document to the set each time you find a new kind of error in real work, rewritten on invented facts. Over time, your test set becomes a record of everything you have learned about how the method can fail, and the fastest way to check whether a tool update has changed anything.

Appendix F: Asking for Budget

Sooner or later, using AI well costs money: a business plan for a tool, a set of tested materials, a training course or time from outside help. Asking finance for budget can feel uncomfortable, especially when the case rests on your own team’s workload. This appendix sets out a short, honest request that finance can approve.

What finance needs to see

Most budget holders need the same things: what you are asking for and what it costs; why now; what it replaces or avoids; what the return is likely to be, and on what basis; what the risks are and how they are controlled; and the decision you want, by when.

Order these for your reader. A finance leader usually wants cost and return first. A general counsel usually wants risk and controls first. A procurement team usually wants the vendor, the terms and the paperwork first.

A one-page request

For a small purchase, one page is usually enough:

Request. Approve [AMOUNT] from [BUDGET LINE] for [WHAT], by [DATE].

Why now. [TWO SENTENCES ON THE WORKLOAD AND WHAT HAS CHANGED, IN YOUR OWN WORDS.]

What changes. [ONE OR TWO SENTENCES ON WHAT THE TEAM WILL BE ABLE TO DO.]

Cost compared with alternatives. [THE COST, AND WHAT THE ALTERNATIVE COSTS, SUCH AS OUTSIDE COUNSEL TIME OR DELAYED WORK.]

Return. [YOUR OWN ESTIMATE, WITH ITS BASIS: FOR EXAMPLE, YOUR TIMED TESTS FROM CHAPTER 10 MULTIPLIED BY YOUR ACTUAL VOLUME.]

Risks and controls. [THE TOOL IS APPROVED ON A BUSINESS PLAN; OUTPUT IS CHECKED UNDER THE TEAM STANDARD; NO PRIVILEGED MATERIAL IS USED.]

Report back. We will report the hours saved, the tasks covered and any issues after [90] days, and recommend whether to renew.

Making the return honest

The return is the part finance will look at most closely, and the part most often overstated. Use your own measurements wherever you can, and say how you made them. If you use a published figure, quote its source and describe it as that source's finding. If your estimate depends on assumptions, such as how many agreements you will review next year, state them.

A modest, well-supported estimate is more persuasive than a large, unsupported one. It also protects you at the 90-day report, because you will be measuring against a figure you can meet.

After approval

Do what the request said. Keep the records, run the 90-day report and recommend honestly whether to renew. A team that reports results, including disappointing ones, will find the next request easier.

Appendix G: For Paralegals, Legal Operations and Legal Assistants

Much of the practical work of using AI in a legal team falls to people who are not lawyers. You may be the one asked to “find ways to use AI”, to build the templates, to run the first passes and to keep the records. This appendix speaks to you directly: what is yours to do, where a lawyer decides and how to build the skills that make you indispensable to the team.

What is yours to do

The methods in this book divide naturally into preparation, first passes, checking and judgment. Preparation is often yours: building test documents, drafting prompts, organizing playbooks and setting up trackers. First passes are often yours too: running the prompt on an incoming agreement, sorting the inbox, drafting the monthly report from the tracker.

Structured checking can be yours, using the “what to check” list for the task: numbers, dates, quotes, missing fields and the specific errors the list names. These checks are careful, skilled work, and doing them well is what makes a first pass safe to pass on.

Where a lawyer decides

Judgment stays with a lawyer: whether a departure matters, what to accept, what to say to a counterparty, what the business should do and anything that might be legal advice. Your team standard should say where these points are, task by task. If it does not, ask your supervising lawyer to mark them, because a clear line protects you as much as it protects the work.

When you are unsure whether something is a judgment call, treat it as one. Flag it, explain what you found and let the lawyer decide. A short note (“Section 9 allows the vendor to change the security standard on notice; I have not marked this as a departure because our positions do not cover it”) is the kind of flag lawyers value.

Building your skills

Three skills make the biggest difference. Writing clear prompts, which is the same skill as writing clear instructions. Checking deliberately, using a list, rather than reading and hoping something stands out. And keeping good records, which turns your experience into evidence the team can use.

The thirty-day plan in Appendix D works as well for you as for a lawyer. Choose a task you already do, test it on invented documents, write the list and keep the log. At the end of the month, you will have something concrete to show.

Making the case to use AI in your role

If you want to use AI in your work and the team has not yet decided, start with the preparation that needs no approval: invented test documents, draft prompts and a written description of the task and its checks. Then ask your supervising lawyer for a short trial on one task, with every output reviewed. Bring timings and an error log to the conversation afterward.

This approach shows the care the team needs to see, and it gives the lawyer something specific to approve.

Questions people ask

What if I am asked to use a tool I think is not approved? Ask, in writing, which tool and plan you should use and for what material. If the answer is unclear, use only invented or public material until it is clear.

What if a lawyer relies on my first pass without checking it? Follow your team standard and record your own checks. If the standard says a lawyer signs off and that is not happening, raise it with the team lead. The standard exists to make the line clear for everyone.

Can I suggest changes to the team's prompts? Yes. Test your change on the team's test documents first and bring the results. The people closest to the work often find the best improvements.

Appendix H: For Small-Business Owners Without In-House Counsel

If you run a small business, you may have no lawyer on staff and no budget to send every contract to an outside firm. You still sign agreements, write policies and make decisions with legal consequences. AI tools can help you understand what you are signing and prepare better questions for a lawyer. They cannot replace a lawyer's judgment, and this appendix is clear about where that line sits.

What AI can help you do

AI can help you read an agreement more quickly and understand it better. Ask it to summarize the agreement in plain terms, with section numbers and quoted wording, using the template in Appendix A. Ask it to list the clauses that limit what you can recover, what you would owe if something went wrong, how the agreement renews and how you can get out of it.

AI can also help you prepare. If you plan to talk to a lawyer, a structured summary and a short list of questions make the conversation shorter and more useful. And AI can help you draft simple documents, such as a first version of a policy, that a lawyer then reviews.

What to check yourself

Use the checking method in Chapter 5. For a small business, five checks matter most:

- **Money.** Every price, fee, increase and penalty, checked against the agreement.
- **Renewal.** Whether the agreement renews automatically, and the date by which you must give notice to stop it.
- **Exit.** How you can end the agreement, and what it costs to do so.
- **Liability.** What you would owe if something went wrong, and whether there is a limit.
- **Your data.** What the other party may do with information about you and your customers, and whether you get it back at the end.

When to call a lawyer

Call a lawyer when the agreement is large relative to your business, long, or hard to exit; when it involves your property, your employees or your customers' personal data; when someone has threatened a claim or sent a demand; when a regulator contacts you; or when the summary leaves you unsure what you are agreeing to. AI can prepare you for that call. It cannot make it unnecessary.

Using AI safely in a small business

The confidentiality questions in Chapter 3 apply to you too. Check the plan you are using and its terms on training and retention, and avoid pasting customer personal data or anything you have promised to keep confidential unless you are confident the tool's terms allow it. Removing names and personal details before pasting is a good habit.

Questions people ask

Can I rely on an AI summary to sign a contract? You can use it to understand the contract, after checking it against the document. Whether to sign is your decision, and for anything significant, a lawyer's review is worth the cost.

Can AI write my contracts? It can produce a first draft, but a draft that looks complete may miss terms that matter for your business or your state. Have a lawyer review any contract you will rely on.

How do I find the right lawyer? Local bar associations often run referral services, and many lawyers offer fixed-fee reviews for small businesses. A clear summary and a list of questions will help any lawyer give you a faster, cheaper answer.

Appendix I: Questions Leaders Ask About Legal's AI Use

Once a legal team starts using AI, leaders will ask about it. Some will want legal to move faster; others will worry that legal is taking risks. These are the questions that come up most often, with the kind of answer a small team can give from the records this book recommends.

"Are we using AI in legal?"

Answer with the team standard. "Yes, for four tasks: NDA review, vendor agreement summaries, contract database summaries and first drafts of board materials. We use the company's approved tool on the business plan, and every output is checked against the source before it is used. A lawyer signs off on anything that goes outside the team."

A specific answer reassures both the leader who wants progress and the one who worries about risk.

"Is our information safe?"

Answer with what you checked and when. "We use the business plan under the company's agreement. Training on our input is off, which IT confirmed on [DATE]. We keep privileged material and unnecessary personal data out of the tool. We recheck the terms every six months and when the provider announces a change."

Avoid absolute assurances. "Safe" is a conclusion; the facts you checked are what the leader can rely on.

“How much time does it save?”

Answer with your measurements, stated with their basis, as Chapter 10 describes. “In our timed tests on [NUMBER] sample NDAs in [MONTH], review took a median of [X] minutes with AI, including checking, against [Y] without. We handle about [VOLUME] NDAs a month, so that is roughly [HOURS] hours a month, which goes to [WHAT THE TEAM DOES WITH IT].”

If you have not measured yet, say so, and say when you will.

“What could go wrong?”

Answer with the errors you have found and how you catch them. “The main risk is output that looks right but misses a condition, such as an exception to a liability cap. Our checklist for each task covers the errors we have seen in testing, and we log what we find. In [PERIOD], our checks caught [NUMBER] errors before anything was used.”

A leader who hears that errors happen and are caught will usually trust the process more than one who hears that errors never happen.

“Why not use AI for more?”

Answer with your method for adding tasks. “We add a task after testing it on invented documents and showing that our checks catch its errors. [TASK] is in testing now. Some tasks, such as [EXAMPLE], depend on judgment the tool cannot supply, so we use it only for the routine parts around them.”

“Should we buy a legal AI product?”

Answer with what you would need it to do and how you would test it. “It depends on whether it does our highest-volume tasks better than the tool we already have, on our documents, with our checks. We would want to test it on our test set, review its data terms and confirm where a person approves its work before deciding.” Chapter 12 lists questions to ask a vendor about agents.

“Are we allowed to do this?”

Answer with the policy and the guidance. “Our use fits the company’s AI policy. Professional guidance, including the ABA’s Formal Opinion 512 and our state bar’s guidance, addresses competence, confidentiality, supervision and communication when lawyers use these tools, and our team standard is designed around those duties.” Check the guidance that applies in your jurisdiction before giving this answer.

Keeping leaders informed

A short paragraph in your monthly report, as Chapter 8 describes, is usually enough: which tasks use AI, any change in tools or terms, any errors caught and any measurement results. Regular, factual updates prevent surprises, and they build the trust that makes the next request easier.

Appendix J: A Sample Note on How Legal Uses AI

This sample note, for an invented company, shows how a legal team might explain its AI use to colleagues in the rest of the business. Adapt it to your team standard and send it once, then link to it from your intake form.

How the legal team uses AI

We use AI to speed up the first steps of some routine legal work, so we can get back to you sooner. The company's approved AI tool helps us prepare first drafts and summaries of NDAs, vendor agreements and some reports. A member of the legal team checks every output against the original document before it is used, and a lawyer makes every legal decision.

We do not put privileged legal advice into the tool, and we remove personal details that a task does not need. We use the tool on the company's business plan, under terms the company has reviewed.

What this means for you

You may get a summary or a first reply from us faster than before. It has been checked, and you can rely on it as you would any note from legal. If something in it looks wrong to you, tell us, because you often know the deal better than we do.

What we ask of you

Please send requests through the legal intake form, with the documents and the deadline. The more complete the request, the faster we can help. Please do not paste our advice, or confidential contracts, into AI tools that the company has not approved.

Questions

Questions about how legal uses AI go to [NAME OR ROLE]. Questions about the company's AI policy go to [POLICY OWNER].

Endnotes

All sources were checked on 26 October 2026.

What did it catch for you? Tell us at provisoiq.com/feedback. Publishing anything you send is your choice.