Lincoln Law

A Hands-On AI Workshop for Attorneys

Worksheets, walkthroughs, and ready-to-build projects

The Workshop Packet

Sixteen worksheets across five tracks, plus companion materials.

Presented by Arthur Rothrock · Legion

June 2026

Contents

Companion Files

In the "Accompanying Materials" folder — separate files, not part of this PDF.

Slide deck — AI Foundations presentation (the deck behind the session).

Sample transcript — Quantum Dynamics PMK deposition of Wong — for worksheet 04.

Sample pitch output — Meridian Litigation Pitch deck — the Claude Design result for worksheet 15.

Research dossier — An AI-built dossier on the SCCBA room — a show-and-tell of AI research.

AI Foundations: The Boring Stuff

What AI Is, What California Requires, and the Policies That Keep You Out of Trouble

---

The rest of this series teaches you to build things with AI. This document is the part you read first, or at least keep in the drawer: what the technology actually is, what California's rules require of you when you use it, how to adopt it at a firm without creating a disciplinary record, and two model policies (one cautious, one permissive) to start from. It is a reference, not legal advice, and the authorities cited here change; verify them before relying on any of this in a matter or before your firm's management committee.

Current as of June 2026. The California State Bar updated its Practical Guidance on May 14, 2026, and amendments weaving AI into the Rules of Professional Conduct are moving through approval. Check calbar.ca.gov for the latest before quoting any of this.

Part 1 — What Is AI, Actually?

"AI" in this document means generative AI: tools like Claude, ChatGPT, Gemini, and Copilot, plus the legal-specific platforms built on the same technology. Under the hood they are large language models, and one fact about them explains nearly every rule in this packet:

A language model does not look things up. It predicts. Trained on enormous amounts of text, the model learned the patterns of language, including legal language, and generates the most plausible next words given what came before. There is no database of cases inside it, no file room it consults. When it writes "(*Smith v. Jones* (2014) 225 Cal.App.4th 1012)", it is producing text shaped like a citation, which may be real (the pattern matches something it learned well) or invented (the pattern is plausible but corresponds to nothing). Both arrive in the same confident prose. Lawyers call the invented ones hallucinations; the model is doing the same thing in both cases.

A few more concepts that matter in practice:

- **Training cutoff.** The model's knowledge stops at its training date. Anything after that (new cases, amended statutes, current addresses) it either doesn't know or guesses. Tools with web search or research databases attached can retrieve current sources, which changes the risk profile: retrieval grounds the output in real documents, though the summary of those documents still needs checking.

- **Context window.** The model only "knows" what is in the current conversation plus its training. Each new chat starts blank, which is why the projects concept from this series (standing instructions plus uploaded reference files) matters: it is how you stop re-explaining yourself.
- **Consumer vs. enterprise vs. legal-specific.** The same underlying model can sit behind a free consumer app (which may use your inputs to train future models), an enterprise deployment (which contractually does not), or a legal research platform (which adds citation databases and retrieval). The ethics analysis below turns heavily on which one you are typing into.
- **Agentic AI.** The newest tools don't just answer; they act: browsing, clicking, running searches, filling forms, executing multi-step tasks. The State Bar's May 2026 guidance update specifically addresses agentic tools. The supervision duty scales with autonomy: the more the tool does on its own, the more actively you must watch it.

The working mental model: a brilliant, tireless first-year who has read everything and verified nothing, writes with total confidence, and has no idea what happened after their last day of reading. Useful beyond measure, and never to be signed off on unreviewed.

Part 2 — Your Ethical Obligations (California)

California has not created an AI exception to the Rules of Professional Conduct, and it has not needed to: the existing rules reach everything that matters. The roadmap is the State Bar's Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law, issued by the Committee on Professional Responsibility and Conduct on November 16, 2023, and updated by the Board of Trustees on May 14, 2026. The guidance is framed as guiding principles rather than binding rules, but two developments give it teeth. First, in August 2025 the California Supreme Court directed the State Bar to consider folding the guidance into the formal rules, and COPRAC has proposed amendments writing AI obligations into the comments of existing rules, where they would become enforceable through discipline. Second, the courts are already enforcing the underlying duties: see the *Noland* discussion below. For attorneys practicing outside California, ABA Formal Opinion 512 (July 29, 2024) reaches parallel conclusions under the Model Rules.

Confidentiality (Rule 1.6; Bus. & Prof. Code § 6068(e))

Inputting client information into an AI tool is a disclosure to the tool's provider. Whether that disclosure is permissible depends on the tool: a consumer product whose terms allow inputs to be used for training or human review is a different animal from an enterprise deployment with contractual confidentiality, encryption, and no-training commitments. The guidance tells you to read the terms of use and understand how inputs are stored, shared, and used before any client information goes in, and to anonymize or withhold what the tool does not need. ABA Opinion 512 goes further for self-learning tools: informed client consent before inputting representation information, and boilerplate consent buried in an engagement letter is not sufficient. Pair your engagement agreement's AI clause with an actual conversation.

Competence and diligence (Rules 1.1, 1.3)

The duty of competence has included technological competence in California since the comment to Rule 1.1 was updated in 2021, and the AI guidance applies it directly: understand the tool's capabilities and limits before using it on client work. Critically, the guidance says competence requires more than detecting and eliminating false AI output; professional judgment cannot be delegated to the tool and remains yours at all times. The guidance also warns against overreliance that atrophies the lawyer's own analysis: AI-generated research should be supplemented by human research and AI-generated argument by human analysis, not the other way around.

Candor to the court (Rules 3.1, 3.3) and *Noland*

Noland v. Land of the Free, L.P. (Sept. 12, 2025) is California's first published appellate opinion on AI fabrications, and it was published, in the court's words, as a warning. Counsel filed appellate briefs in which nearly all the quoted authorities were invented by generative AI: passages attributed to opinions that never contained them, cases that did not address the issues claimed, and authorities that did not exist. The Court of Appeal imposed $10,000 in sanctions, ordered the opinion served on the client, and referred counsel to the State Bar. The rule the court stated is absolute and applies to every source, AI or otherwise: no brief, pleading, or motion should contain any citation the responsible attorney has not personally read and verified. One more detail worth absorbing: the court declined to award sanctions to the respondent partly because respondent's counsel never noticed the fabrications either. Verification is everyone's job, including when reviewing the other side's papers.

Supervision (Rules 5.1, 5.2, 5.3)

Managerial lawyers must establish clear policies on permissible AI use and take reasonable measures to ensure lawyers and staff comply, including training on the ethical and practical pitfalls. That is the direct source of the model policies in Part 5: under the guidance, having a policy is not optional once your people are using these tools, and they are using them whether or not you have asked. The guidance also addresses subordinates: a junior lawyer cannot use AI improperly even at a supervisor's direction.

Communication (Rule 1.4)

Consider whether your AI use is something the client would reasonably want to know about, and disclose where it is. Disclosure is clearly indicated where AI will process the client's confidential information, where its costs will be charged, or where the client has asked. The engagement agreement clauses that accompany this packet (Addendum A on the firm's use, Addendum B on the client's) are the written half of that communication; the conversation is the other half.

Fees (Rule 1.5)

Bill for the time you actually spend, including time spent crafting prompts and reviewing output, and not for the time the tool saved you. If a brief that once took ten hours now takes three, you bill three. The guidance contemplates passing through tool costs where disclosed and agreed, which is what Addendum A's cost provisions do. ABA Opinion 512 adds detail: general AI subscriptions are typically overhead, you may not charge clients for teaching yourself a tool you will use across your practice, and any per-use cost passed to a client must be explained in advance.

Bias and discrimination (Rule 8.4.1)

Models learn from human text and inherit its biases. The guidance flags this for any AI use touching evaluative judgments: screening clients, employment decisions, assessing witnesses or jurors. Know that the skew exists, and keep a human decision between the model's output and any consequential call about a person.

The courts themselves (Rule of Court 10.430)

Effective September 2025, California adopted Rule of Court 10.430 and Standard 10.80, requiring courts that permit generative AI use by judicial officers and staff to adopt policies on confidentiality, privacy, bias, safety, supervision, and transparency. The rule governs the courts' own use, not attorney filings, but it tells you where the judiciary stands: AI is not prohibited; unsupervised AI is unacceptable. Separately, individual judges increasingly impose AI disclosure or certification requirements through standing orders. Add "check the department's standing orders for AI provisions" to your filing checklist the same way you check formatting rules.

Part 3 — Adopting AI: Best Practices for the Firm

Adoption is a governance project, not a software purchase. The sequence that works:

1. **Name an owner.** One attorney (with IT support if you have it) owns the AI policy, the approved-tool list, and the questions. Committees can advise; one person decides and updates.
2. **Find the shadow AI first.** Survey, without punishment attached, what people are already using. The answer is more than you think, on personal accounts, which is the riskiest possible configuration. Policy should channel existing behavior, not pretend it doesn't exist.
3. **Vet tools before approving them.** The vendor diligence questions below are the minimum. Approve specific tools at specific tiers (see the model policies), and revisit the list quarterly; this market changes monthly.
4. **Put it in writing.** Adopt a policy (Part 5 gives you two starting points), have everyone acknowledge it, and keep the acknowledgments. Under Rules 5.1 and 5.3, the policy plus training is how managerial lawyers discharge the supervision duty.
5. **Train before, not after.** The guidance expressly includes training in the supervision duty. One hour covering hallucinations, confidentiality tiers, verification, and billing honesty prevents most of the failure modes. The worksheets in this series are the curriculum.
6. **Tell the clients.** Update the engagement agreement (the companion clauses do this), and have the conversation with clients whose matters will involve AI processing of their information. Some institutional clients have their own outside counsel guidelines restricting AI use; check them before assuming.
7. **Document and revisit.** Keep a simple record of what tools were approved, when, and on what diligence. Revisit the policy on a calendar, not when something breaks.

Vendor diligence questions (ask before approving any tool)

```
1. Are our inputs used to train your models, ever, under our tier? Where is that
commitment in the contract?

2. Who can see our data: your employees, human reviewers, subprocessors? Under
what circumstances?

3. How long are prompts and outputs retained, and can we set retention to zero or
near it?

4. Where is the data stored and processed, and is it encrypted in transit and at
rest?

5. What security certifications do you hold (e.g., SOC 2 Type II), and can we see
the report?
```

```
6. Will you sign confidentiality and data-protection terms consistent with our obligations to clients?

7. What happens to our data when we leave?

8. For legal research tools: where do your citations come from, and does the tool retrieve real sources or generate text?

9. For agentic tools: what actions can it take without human approval, and what are the controls?
```

Part 4 — Using AI: Best Practices Day to Day

- **Match the data to the tool.** Before typing, ask: what tier is this tool, and what tier is this information? Client confidences belong only in tools approved for them. When in doubt, anonymize: initials, roles, rounded numbers.
- **Verify on a ladder.** Brainstorming needs a sniff test. A client email needs a careful read. Anything filed needs every citation pulled and read, every quotation checked against the source, by the attorney signing it. *Noland* is the floor, not the ceiling.
- **Write prompts like assignments.** The CASE structure from the FRAME worksheet (Context, Ask, Specs, Examples) is the difference between output you can use and output you rewrite. Include "do not invent facts, cases, or quotations; flag uncertainty" in anything substantive, and remember that the instruction reduces the risk without eliminating it.
- **Keep the chat out of the file, but keep a record.** Know your tool's retention settings. For significant work, note in the matter file what was AI-assisted and what verification was done; if a court or carrier ever asks, the answer should be on paper.
- **Bill the truth.** Actual time, including prompting and review. Disclosed tool costs as agreed. Nothing for the hours the machine saved you.
- **Check the courtroom.** Before filing, check the judge's standing orders and local rules for AI disclosure or certification requirements, which vary by department and change.
- **Watch the agent.** For browser-driving and other agentic tools: your logins, your approvals, supervised sessions, and no standing permissions on court, billing, or banking sites.
- **Escalate surprises.** When AI output materially changes your view of a question that matters, that is a conversation with a colleague, not just another prompt.

Part 5 — Two Model Firm Policies

Two starting points, written to be edited, and informed by the structure of widely used standard policy documents. The conservative version suits firms that want control first and expansion later; the permissive version suits firms that want adoption with guardrails. Both share the same non-negotiable floor: verification, confidentiality tiers, billing honesty, and candor. Neither is legal advice; adapt to your practice, your clients, and your malpractice carrier's views, and see the flags at the end of this Part before finalizing.

Model Policy A: The Conservative Firm

```
FIRM POLICY ON ARTIFICIAL INTELLIGENCE (RESTRICTIVE)

1. Scope. This policy applies to all attorneys and staff in the course of firm or
client work, on any device (firm or personal), and extends to AI features
embedded in other software (word processors, email, research platforms). It
supplements, and does not replace, the firm's confidentiality, information
security, and other policies.

2. Approved tools only. Personnel may use only the AI tools on the firm's
Approved Tools List, accessed exclusively through firm-administered accounts and
firm credentials. Use of personal accounts or unlisted tools for any firm or
client work is prohibited. Access ends when employment ends.

3. Authorized purposes. Approved tools may be used for: research from public
sources, first drafts from attorney-provided material, summarization, and
internal templates and checklists. Any other use, and any use of AI in hiring,
personnel evaluation, or client-screening decisions, requires the managing
partner's advance written approval.

4. Client information. No client-identifying information may be entered into any
AI tool without the prior written approval of the supervising partner for that
matter. Where approval is given, information must be limited to what the task
requires and anonymized where feasible. No information subject to a protective
order, NDA designation, or heightened confidentiality (medical, minors, trade
secrets) may be entered into any AI tool absent the managing partner's written
approval.

5. Verification. All AI output is treated as an unverified draft. Every factual
statement, citation, and quotation must be verified against primary sources by an
attorney before use. No AI-assisted document may be filed with any court or sent
outside the firm unless all facts and authorities have been independently
verified; filings additionally require a second attorney's cite check.

6. Records. Personnel log AI use on client matters in the matter file (tool,
task, verification performed). The firm may monitor use of firm-administered AI
accounts, and may audit inputs and outputs for compliance.

7. Client disclosure. The firm's engagement agreement discloses our use of AI.
For any matter in which client confidential information will be processed by an
AI tool, the responsible attorney will obtain the client's informed consent and
document it in the file. Client outside-counsel guidelines restricting AI use
control over this policy.
```

```
8. Billing. Bill only actual time spent, including prompt drafting and output
review. Disclosed third-party AI costs may be passed through as costs where the
engagement agreement provides. Time saved through AI may not be billed.

9. Court rules. Before any filing, the responsible attorney confirms compliance
with all applicable standing orders and local rules on AI disclosure.

10. Vendors and co-counsel. Personnel supervising contract attorneys, vendors, or
co-counsel must confirm their AI use on firm matters complies with this policy or
equivalent safeguards.

11. Training. All personnel complete AI training before access and annually
thereafter; the firm retains training records. Managers ensure completion.

12. Reporting and violations. Suspected violations, and any unapproved AI tool
found on firm systems, must be reported promptly to the managing partner. The
firm prohibits retaliation for good-faith reports. Violations may result in
discipline up to termination; prompt self-reporting is a mitigating factor,
concealment is not.

Owner: [NAME]. Effective date: [DATE]. Review: [QUARTERLY]. Revision history:
[DATE / DESCRIPTION].
```

Model Policy B: The Permissive Firm

FIRM POLICY ON ARTIFICIAL INTELLIGENCE (PERMISSIVE)

1. Scope and intent. The firm encourages the use of AI to improve quality and efficiency. This policy applies to all attorneys and staff doing firm or client work on any device, and covers standalone AI tools and AI features embedded in other software. It supplements the firm's confidentiality, security, and personnel policies.

2. Tiers. Tier 1 (firm enterprise tools with contractual no-training and confidentiality commitments): approved for client work, including confidential information, subject to this policy. Tier 2 (reputable consumer tools on firm-administered accounts): approved for non-confidential work only: research on public law, drafting from sanitized facts, internal templates, marketing. All access through firm credentials; access ends at separation.

3. The confidentiality line. Client-identifying information and client documents go into Tier 1 tools only. Protective-order material, sealed records, and information the client has restricted go into no AI tool without the supervising partner's written approval. When in doubt, sanitize or ask.

4. The verification floor (non-negotiable). Every citation, quotation, and factual assertion in work product is verified against primary sources by the responsible attorney. Court filings: the signing attorney personally confirms every cited authority, per *Noland v. Land of the Free*. Nothing AI-assisted leaves the firm without attorney review. No exceptions for any tool, including legal research platforms.

5. Off-limits uses. No AI use for hiring, personnel evaluation, or client-screening decisions without the AI lead's written approval and bias review; no uses that are illegal, infringing, or harassing.

```
6. Experimentation. Personnel are encouraged to pilot new tools on non-client
work and propose them for the Tier lists. The AI lead maintains the lists, runs
vendor diligence, and updates classifications monthly. Share what works; the firm
maintains a prompt-and-project library.

7. Client communication. The engagement agreement discloses AI use. The
responsible attorney exercises judgment under Rule 1.4 about further disclosure,
and obtains informed consent before a client's confidential information is
processed where required. Client outside-counsel guidelines restricting AI
control over firm preference.

8. Billing. Actual time only, including prompting and review. Disclosed tool
costs as the engagement agreement provides. Efficiency gains are passed to the
client or addressed through alternative fee arrangements, never through phantom
hours.

9. Vendors, co-counsel, and agents. AI use by contract attorneys, vendors, and
co-counsel on firm matters must meet this policy's confidentiality and
verification standards; supervising attorneys confirm it. Agentic tools run only
in supervised sessions: user-held credentials, action-by-action approval on court
and financial sites, session logs retained.

10. Records. Significant AI assistance on client work is noted in the matter file
with the verification performed. The firm may review use of firm-administered AI
accounts for compliance.

11. Reporting, training, and culture. AI training at onboarding and annually.
Report unapproved tools or misuse to the AI lead; no retaliation for good-faith
reports. Mistakes surfaced promptly are treated as learning; concealment is
treated as a violation.

Owner: [NAME]. Effective date: [DATE]. Review: [MONTHLY FOR TIER LISTS, ANNUALLY
FOR POLICY]. Revision history: [DATE / DESCRIPTION].
```

Flags to consider before you adopt either policy

- **Personal devices and personal accounts.** The riskiest AI use happens on a personal phone with a personal login, outside every safeguard. Both policies now reach it by scope; decide how you will actually detect and address it, and whether your IT setup supports firm-credentialed access.
- **Embedded AI is the new unseen gap.** AI features now ship inside word processors, email, research platforms, and practice management software. A policy that only names chatbots misses most of the surface. Inventory the AI features in software you already license and classify them like any other tool.
- **Monitoring and employment law.** If the firm reserves the right to monitor AI use, run the provision past employment counsel: monitoring and discipline policies must account for the National Labor Relations Act (which applies to non-union workplaces too) and, in California, employee privacy law. The anti-retaliation language in both policies is there for a reason; keep it.

- **Records cut both ways.** Usage logs and saved chats help with supervision, audits, and proving verification, and they are also discoverable if not work product. Decide deliberately what gets logged, where, and for how long, and align it with your retention schedule rather than defaulting to keep-everything or keep-nothing.
- **HR and hiring uses carry their own regime.** AI in employment decisions implicates anti-discrimination law and, in California, specific regulations on automated decision systems. Both policies carve these uses out for separate approval; if your firm wants to use AI in hiring, that is its own compliance project.
- **Flow-down to vendors and co-counsel.** Your obligations don't stop at your own keyboards. Contract attorneys, e-discovery vendors, and co-counsel using AI on your matters can create the same exposures; both policies now require supervising attorneys to confirm equivalent safeguards. On the receiving end, check institutional clients' outside counsel guidelines, which increasingly include AI restrictions that control over your policy.
- **The regulatory ground is moving.** State AI laws, the State Bar's pending rule amendments, and federal efforts to preempt state AI regulation are all in motion as of this writing. The effective-date and revision-history block at the bottom of each policy is not decoration; calendar the review and record what changed and why.

Part 6 — Security and Privacy

The training question comes first

The single most consequential setting in any AI tool: are your inputs used to train future models? Consumer tiers often say yes by default; business and enterprise tiers often contractually say no. For client work, the answer must be no, in the contract, not in a marketing page.

Retention, access, and accounts

- **Retention settings.** Know how long chats are stored and set retention as short as the tool allows for sensitive work.
- **Firm-administered accounts.** All work happens on firm accounts with SSO or at minimum unique credentials and multi-factor authentication. Personal accounts are how confidential info end up outside every safeguard the firm has.
- **Least privilege for integrations.** Tools connected to email, documents, or calendars see everything they are connected to. Scope the connections to what the use case needs.

Vendors

Hold AI vendors to the standard you hold any vendor holding client confidences: security certifications you can inspect, encryption in transit and at rest, breach notification commitments, subprocessor transparency, and contractual confidentiality. The diligence questions in Part 3 are the script.

The threat side

- **Deepfakes and voice cloning.** Wire instructions, "client" calls with changed payment details, and urgent requests from "the managing partner" can now be convincingly faked. Verify out-of-band, by a number you already had, before money moves.
- **Prompt injection.** Agentic tools reading web pages or documents can be manipulated by instructions hidden in that content. Supervise sessions and treat unexpected behavior as a stop signal.
- **Incident response.** If confidential information goes into the wrong tool, treat it like any potential disclosure: assess, contain (most vendors can delete on request), evaluate notification duties to the client, and document. Add AI mishaps to your existing incident plan, and ask your malpractice carrier how your AI use affects coverage before you need the answer.

Part 7 — Engagement Agreement Clauses (Copy-Paste)

Two ready-to-use clauses. Lift them into your engagement agreement or attach as an addendum, and edit to fit your practice. Not legal advice. Remember the lesson from ABA Opinion 512: where informed consent is required, the signed clause is necessary but not sufficient; pair it with an actual conversation, and document that the conversation happened.

Addendum A: The Firm's Use of AI

```
USE OF ARTIFICIAL INTELLIGENCE TOOLS

To assist in providing efficient and cost-effective legal services, our firm may
use artificial intelligence (AI) platforms when working on your matter, including
tools designed specifically for legal professionals that assist with drafting,
research, and document review.

We use AI platforms that maintain appropriate confidentiality and security
protections, including encryption and data isolation. We do not input
confidential client information into public AI tools that lack adequate
safeguards. We will review all AI-generated output for accuracy and completeness,
and we will independently verify any legal authorities before relying on them in
your matter.

We will bill only for the actual time we spend using AI tools and reviewing their
output, not for time hypothetically saved. Some AI tools carry third-party costs,
which we will disclose to you and may charge to your matter as a cost.

By signing this agreement, you acknowledge that you have read and understand the
above and consent to our firm's use of AI tools as described.
```

Example cost-disclosure paragraph for per-page AI tools, such as Legion: add a paragraph like the following when a specific tool's charges will be passed through as costs.

```
Our firm uses Legion, an AI drafting platform built for litigators that bills on
a per-page basis. Because Legion can complete in hours work that would
traditionally take days or weeks, it often produces significant cost savings and
better results for you. These charges are billed to your matter as a cost and
itemized on a monthly invoice we will provide to you.
```

Addendum B: The Client's Use of AI

```
USE OF ARTIFICIAL INTELLIGENCE

"AI tools" means chatbots and other programs that generate, summarize, or rewrite
text for you, such as ChatGPT, Claude, Gemini, and Copilot, and the AI features
now built into search engines, email, and word processors.

Why this matters for your case.
```

Your private information can stop being private. Normally, what you share with us is protected by something called the attorney-client privilege, a rule that keeps our communications confidential and out of the other side's hands. If you type details about your case into a public AI tool or Google search, you are handing that information to an outside company, and that can destroy the protection. Once it is gone, the other side may be entitled to demand and read what you shared. (Lawyers call information the other side can force you to turn over "discoverable.")

The AI can simply be wrong. These tools regularly produce answers that are incomplete, misleading, or entirely made up, including fake court cases that look real. Relying on them can give you a false picture of your case and cost extra time and money to untangle.

It can get you in trouble with the court. Using AI in your case can violate court rules and orders, which can lead to penalties.

What we ask of you.

Keep case information out of AI tools. Don't enter documents, our communications, or anything you've received in the case.

Don't present AI's work as your own. If you used AI to write or research something you send us, tell us.

Tell us if you've used AI at all in connection with your case, so we can address any problems early.

Rely on us, not the AI, for legal advice. We are your lawyers; the AI is not, and it should not guide decisions about your matter.

How we use AI. We may use AI tools ourselves, for research, drafting, and review, but always with the safeguards described above, and always with a licensed attorney checking the work before we rely on it or send it to you.

If these rules aren't followed. You could lose the privacy protections on your information, face penalties from the court, run up extra fees to fix the problem, and, in serious or repeated cases, we may have to stop representing you.

Client Signature: ________________________________   Date: ____________

Client Name (Printed): ________________________________

Addendum A operationalizes the disclosure and billing duties from Part 2 (Rules 1.4 and 1.5, and the guidance's bill-actual-time principle). Addendum B addresses the half of the problem most firm policies miss: the client's own AI use, which can waive protections and inject fabricated law into your case from the outside. The remaining worksheets in this series cover the affirmative side: the FRAME method, projects, and practice-specific suites.

The One-Page Version

☐ The model predicts; it does not look up. Verify everything against primary sources, always.

☐ No AI exception to the rules: confidentiality (1.6), competence (1.1), supervision (5.1-5.3), communication (1.4), fees (1.5), candor (3.1, 3.3) all apply in full.

☐ Client information goes only into tools with contractual no-training and confidentiality commitments, and informed consent comes before confidential information goes in where required.

☐ Every filed citation personally read and verified by the signing attorney. *Noland*: $10,000, State Bar referral, published as a warning.

☐ A written policy, training, and an approved-tool list are how managers discharge the supervision duty.

☐ Bill actual time and disclosed costs. Never time saved.

☐ Check standing orders before filing; the disclosure landscape is department by department.

☐ Secure tools protect confidentiality, not privilege. Write prompts accordingly.

The point of the boring stuff: every rule in this packet reduces to one sentence the profession has always lived by: the lawyer is responsible for the work. The tools changed; the sentence didn't.

Stop Re-Explaining Yourself to AI

Legion's FRAME Method and Three Ready-to-Build AI Projects for Attorneys

FRAME is an original framework I developed at Legion. It isn't borrowed from any course or prompt library.

Every time you open a fresh AI chat, the assistant knows nothing about you: not your practice area, not how you write, not what you consider good work product. So you re-explain it. Every time. This worksheet fixes that in two moves. First, a working method, FRAME, for getting AI to do anything you need. Second, "projects": persistent workspaces in Claude, ChatGPT, or Gemini that carry your standing instructions and reference files into every conversation, so the AI already knows who you are before you type a word. You'll set up three of them today.

Part 1 — Legion's FRAME Method

You don't need prompt libraries or a "prompt engineering" course. You need a habit. When a task makes you groan, frame it for the AI and let it do the work, then save the frame so you never build it twice.

F — Friction is your trigger

The moment you don't want to do a task is the moment to hand it to AI. Don't overthink whether it's an "AI task." Dread is the signal.

R — Reach out and ask it how

Don't know how to make AI do something? Ask the AI. "How can you help me do this, and what do you need from me?" Let it interview you. It will tell you what to paste in, what it needs to know, and how to break the task up. You are not expected to arrive knowing the answer; curiosity is the skill.

A — Ask well, the CASE way

A good prompt has four parts, plus a fifth when research matters:

- **C — Context:** what you're working on, the documents, the strategy in your head. Find the Goldilocks zone: too much = bad results, too little = bad results.
- **A — Ask:** the task, as one clear instruction. Not "pls fix."

- **S — Specs:** length, tone, format, "don't make anything up," and "here is how I would do it normally: [explain your process]."
- **E — Examples:** paste or upload something that nails the voice or format you want.
- **+ S — Search:** when the task needs it, add: "find the best way to do this, research the key authorities, then do it."

M — Make it better

When the output isn't exactly right, don't start over. Tell it exactly what was wrong. Go another round. Get stuck? Start a new chat.

E — Embed it

Once it's good, save the frame so you never rebuild it. Set up a project that knows who you are and what good output looks like (your brief format, your letter voice, your standing instructions) so you stop re-explaining yourself. Then ask the AI for a reusable prompt after each task: "Write me a prompt I can paste next time to do this exact task." Create a new project, paste in the prompt, save. Now you can stop re-explaining the same thing over and over.

The rest of this worksheet is the E step in action. Parts 2 through 4 turn "embed it" into three concrete projects you build today.

Part 2 — What a "Project" Is and How to Make One

All three major platforms offer the same core idea under different names: a workspace with (1) standing instructions the AI reads at the start of every chat, and (2) a set of reference files it can draw on. Set it up once; every conversation inside it starts fully briefed.

	Claude (Anthropic)	ChatGPT (OpenAI)	Gemini (Google)
What it's called	Projects	Projects (also Custom GPTs for shareable assistants)	Gems
Where to find it	"Projects" in the left sidebar → + New Project	"Projects" in the left sidebar → New project	"Gems" / Gem manager in the sidebar → New Gem
Standing instructions	"Instructions" panel on the project page	Project instructions field	The Gem's "Instructions" box
Reference files	"Files" / project knowledge: upload documents Claude uses in every chat	Add files to the project	Attach files, including directly from Google Drive
Good to know	Chats inside the project all start with your instructions and files; chats don't share memory with each other unless saved to project knowledge	Keeps related chats, files, and instructions in one workspace	Drive-connected files stay current when the source document changes

Setup, in five steps (same on every platform)

1. **Create the project/Gem** and name it for the job, not the tool: "CA Briefs — Drafting," not "Legal AI Helper."
2. **Paste your instructions.** Start with the foundation block from Part 3, then the job-specific instructions from Part 4.
3. **Upload reference files**: your best work product, templates, and anything the AI should imitate or consult.
4. **Test it with three real tasks** you've already done, and compare the output to what you actually produced.
5. **Refine the instructions** based on what missed. A shortcut: paste the bad output back in and say, "Rewrite your own project instructions so you would have gotten this right the first time."

Before you upload anything: the confidentiality gate

☐ I am using a firm-approved AI account (typically a business/enterprise tier where inputs are not used to train models), not a personal free account.

☐ Nothing I upload to project knowledge contains client confidences or privileged material unless the firm's deployment and my engagement terms permit it.

☐ Where client documents are needed, I anonymize or use my own past work product as the example instead.

☐ I have checked my state bar's guidance on generative AI use (California attorneys: the State Bar's Practical Guidance on generative AI).

Part 3 — Your Foundation Block

Every project you build should open with the same short block about you and your firm. Draft it once here, then paste it at the top of every project's instructions. Keep it under half a page; the Goldilocks zone applies.

```
About me: I'm [NAME], a [PARTNER/ASSOCIATE/SOLO] at [FIRM NAME], a [SIZE]-
attorney firm in [CITY/REGION].

I practice [PRACTICE AREAS] in [STATE], primarily in [COURTS/VENUES/AGENCIES].

My typical clients are [TYPES OF CLIENTS] and my typical matters are [TYPES OF
MATTERS].

My standing preferences: [E.G., PLAIN ENGLISH, NO HEDGING, FLAG UNCERTAINTY
EXPLICITLY, ASK ME QUESTIONS BEFORE DRAFTING IF MY REQUEST IS AMBIGUOUS].

Never invent facts, cases, statutes, or quotations. If you are unsure, say so and
ask.
```

Draft yours

Name, role, firm, size, location:

____

Practice areas, state, courts/venues:

____

Typical clients and matters: ____

Standing preferences (how you like to work):

____

Anything the AI should always do (or never do):

____

Part 4 — Three Projects to Build Today

These are real, field-tested project instructions. Paste the foundation block first, then the instructions below. Adapt the bracketed parts and anything that doesn't match your practice.

Project 1 — Drafting: California Briefs

Suggested name: "CA Briefs — Drafting." **Upload to project knowledge:** one or two briefs you're proud of, your caption/formatting template, and the operative documents for the current matter as you go.

```
You help me draft California legal briefs.

No bullet or numbered lists, briefs don't use them. If you must list, do it
inline: (1) ...; (2) ...; and (3) ....

Use Roman numerals for major sections (e.g., Statement of Facts), then A./B./C.
and 1./2./3. only where a section genuinely needs breaking down. Make headings
explanatory and quietly persuasive.

Argue persuasively but never conclusory or over the top: connect the facts to the
law, cut flowery language and snark, no hyperbole.

Follow CREAC (conclusion, rule, explanation, analysis, conclusion) without
announcing it, walking the court to the conclusion.

Italicize case names and Id.; cite after the sentence like this: (Thorley v.
Superior Court (1978) 78 Cal.App.3d 900, 907–908.) Don't italicize statutes.

Never use em dashes.

Never cite a case or statute I haven't given you.

In a facts section, cite the declaration or document for every fact.

Avoid typical AI tells such as "It's not just x, it's y."
```

Project 2 — Correspondence: Letters in My Voice

Suggested name: "Correspondence — My Voice." **Upload to project knowledge:** two or three letters you're proud of, plus the voice-quote file described below.

```
You draft my correspondence, including meet-and-confer letters.

Write the way I do: direct, confident, plain English. Lead with the point.

Short paragraphs, contractions, no corporate throat-clearing ("at your earliest
convenience," "we would be most grateful").

Firm but never hostile or snarky.

Every word earns its place, less is more.

Never write filler like "the thing nobody's talking about" or "it's not just X,
it's Y."
```

To teach it your voice: paste in two or three letters you're proud of and say, "Pull about 12 short quotes that capture how I write, across different situations." Save those quotes as a file in the project's knowledge. Now every letter starts from your actual voice, not a generic lawyer voice.

Project 3 — Strategy: Your Straight-Talk Advisor

Suggested name: "Practice Advisor." **Upload to project knowledge:** a one-page profile of your practice (staffing, rates, typical clients, revenue goals, current headaches) and anything else needed to understand your firm. The more it knows, the less generic the advice.

```
You're a blunt, no-BS advisor for running my law practice.

Tell me when I'm making the right move and when I'm screwing up, don't flatter me.

I want actionable guidance on handling clients, managing staff, pricing and collections, business development, and the hard calls.

[Provide all the details about your practice, employees, typical clients, and any other relevant materials for understanding your firm, either here or as uploaded files.]
```

Part 5 — Build Your Own (Exercise)

Now run FRAME end to end on a task of your own.

1. **F:** Write down the recurring task you dread most: ____________________________
2. **R:** Open a chat and ask: "How can you help me do this, and what do you need from me?" Note what it asks for.
3. **A:** Run the task once using CASE — context, ask, specs, examples (plus search if needed).
4. **M:** Iterate until the output is something you'd actually use. Tell it precisely what's wrong each round.
5. **E:** Ask: "Write me a prompt I can paste next time to do this exact task." Create a new project, paste your foundation block plus that prompt into the instructions, upload your examples, and save.

Project planning fields

Project name: ____________________________

The job this project does: ____________________________

Standing instructions (beyond the foundation block):

____

Files to upload as knowledge: ____________________________

First three test tasks: ____________________________

Quality checklist before you rely on it

☐ Instructions are specific enough that a new associate could follow them. "Write well" is not an instruction.

☐ At least one example of excellent output is in the project knowledge.

☐ The instructions include "never invent facts, cases, or quotations" and tell the AI what to do when uncertain.

☐ Tested against three real tasks and refined at least once.

☐ Confidentiality gate (Part 2) cleared for every file uploaded.

☐ You, not the AI, remain the final reviewer of everything that leaves your desk.

The payoff: a project is an associate who never forgets your instructions and is available at midnight before the filing deadline. Build one per recurring job, and the time you spend explaining yourself to AI drops to zero.

The Litigation Sparring Suite

Four AI Projects That Pressure-Test Your Case, Built with the FRAME Method

---

The projects in this worksheet don't write your work product. They attack it. Each one casts the AI in an adversarial role (opposing counsel, a skeptical judge, an evasive deponent, a witness fighting your cross) so you find the weaknesses in a conference room instead of a courtroom. Nobody on your own team will ever hit your argument as hard as the other side will; these projects make sure something does.

Each project below is built using the FRAME method from the companion worksheet, so you can see the method working four times in a row: Friction tells you when to reach for the project, Reach out shows you the first question to ask, Ask gives you the project instructions and a model CASE prompt, Make it better gives you the iteration moves that matter for that project, and Embed tells you exactly what to name it and what to upload.

FRAME in one breath: **F**riction is your trigger. **R**each out and ask the AI how. **A**sk well, the CASE way (Context, Ask, Specs, Examples, plus Search when research matters). **M**ake it better through rounds, not restarts. **E**mbed it as a project so you never rebuild it.

Before anything else: the confidentiality gate from the Projects worksheet applies with extra force here, because these projects run on real case materials. Firm-approved account, no privileged or sealed material unless your deployment permits it, and check protective orders before uploading transcripts or produced documents.

Project 1 — The Devil's Advocate

Your opposing counsel, on demand. It reads your brief or theory the way the other side will: hunting for the gap in the record, the counter-authority, the alternative reading of your best case.

F Friction is your trigger

The moment you finish a draft and feel good about it. Comfort is the danger sign; by the time you've written something, you've talked yourself into it. Use this project before a brief is filed or a demand goes out.

R Reach out and ask it how

```
I'm about to have you attack my draft motion. Before you do, tell me: what do you
need from me to attack it the way real opposing counsel would? The opposition's
papers? The key record documents? Their counsel's prior briefs?
```

A Ask well: the project instructions

Paste your foundation block first, then these instructions:

```
You are opposing counsel: a skilled, aggressive litigator on the other side of my
case.

When I give you a brief, argument, demand, or case theory, attack it. Find every
factual weakness, every gap in the record, every counter-authority among the
cases I have provided, every alternative reading of the cases I rely on, and
every procedural vulnerability.

Identify the motion, opposition, or response you would file against me and
outline it.

Rank the weaknesses you find from most dangerous to least, and explain what each
one costs me if I don't fix it.

Do not soften the attack to spare my feelings. Do not invent facts, cases, or
record cites: work from what I give you, and where you would need the record to
land a punch, say exactly what document you would demand.

After the attack, switch sides for one section titled "If I Were You": the three
fixes that would hurt you most as my opponent.
```

A model first prompt, the CASE way:

```
CONTEXT: Attached are my draft MSJ, the operative complaint, and the key
deposition excerpts. We represent the employer in a wrongful termination case;
our theory is legitimate business reasons plus documented performance issues.

ASK: Attack this motion as plaintiff's counsel.

SPECS: Rank weaknesses inline as (1), (2), (3). Quote the exact sentence of my
brief you're attacking before each strike. Use only the authorities in my brief
and the attached opposition cases. Don't invent anything.
```

```
EXAMPLES: Attached is plaintiff's counsel's opposition from the Ramirez matter so
you can attack in their actual style.
```

M Make it better

- "Weakness (2) is the scary one. Now write the strongest two paragraphs the opposition would write on it, full strength, no summary."
- "You're attacking generic positions, not my record. Here are the cites; go again and anchor every strike to a page."
- "Now grade the attack: which of your strikes would actually survive my reply, and which are noise?"

E Embed it

Name it: "Devil's Advocate" (or one per major case: "Devil's Advocate — Ramirez"). **Upload:** the operative pleadings, key rulings, your draft, and opposing counsel's prior briefs if you have them; it will learn how your actual opponent argues. Then ask: "Write me a reusable prompt for attacking any brief I paste in," and save it to the project.

Project 2 — The Judge

A skeptical bench officer with a heavy calendar who tells you how the motion is landing, then asks you the questions you'll actually get at the hearing.

F Friction is your trigger

A hearing is on calendar and you're rehearsing your argument in the shower instead of against resistance. The dread of "what will the court push on?" is the signal. Run this after Devil's Advocate, once the brief is filed and the fight is about the hearing.

R Reach out and ask it how

```
You're going to play the judge hearing my motion. What would help you predict the
bench's questions: the tentative ruling practice in this department, the
opposition papers, my judge's published rulings, the local rules?
```

A Ask well: the project instructions

```
You are a skeptical California superior court judge with a heavy law-and-motion
calendar. You have read the papers, but you have limited time and no patience for
overstatement.

When I give you a motion and opposition, do four things:

First, in two paragraphs, tell me how you are tentatively inclined to rule and
why, including what would change your mind.

Second, list the 12 to 15 questions you would ask from the bench, ordered from
most likely to least likely. Include the hostile questions, the questions that
expose my weakest link, the practical questions (what order am I actually asking
you to sign?), and at least two questions directed at my opponent that I should
be ready to capitalize on.

Third, for each question, add one line on what answer would satisfy you and one
line on what answer would sink me.

Fourth, tell me what in the papers annoyed you: overstatement, missing record
cites, authority the parties ignored, requests for relief the notice doesn't
support.

Stay in character. Do not invent cases or facts. Work only from the papers and
authorities I provide.
```

A model first prompt:

```
CONTEXT: Attached are my motion to compel further responses, the opposition, my
reply, and the separate statement. Hearing is in Department 6; this judge issues
tentatives and adopts them about 80% of the time.

ASK: Give me your tentative, your bench questions, and your critique, per your
instructions.
```

```
SPECS: Keep the tentative to two paragraphs. Questions numbered, most likely
first. Be blunt in the "what annoyed you" section.

EXAMPLES: Attached are two of this judge's published tentatives on discovery
motions so you can match the court's reasoning style.
```

M Make it better

- "Now ask me your questions one at a time. After each answer I give, grade it and tell me the follow-up you'd ask."
- "Rule against me. Write the minute order denying the motion so I can see the reasoning I have to cut off."
- "Which three questions should I open my argument by answering before you ask them?"

E Embed it

Name it: "The Judge." **Upload:** the motion, opposition, and reply for the current hearing; the key authorities; the department's local rules; and any published tentatives or rulings from your actual judge. Refresh the files each hearing; the instructions never change.

Project 3 — The Deposition Prep Partner

Two tools in one project: an outline builder that organizes the examination around the admissions you need, and a role-play witness who is every bit as evasive as the real one will be.

F Friction is your trigger

The deposition is next week, the documents are in a banker's box (or a 4 GB production folder), and the outline is a blank page. Blank-page dread plus document overload is exactly the friction this project eats.

R Reach out and ask it how

```
I'm deposing the defendant's HR director in three weeks. Before we build
anything, interview me: ask me what I need from this deposition, what the witness
knows, and which documents matter. Then tell me how you'd structure the prep.
```

A Ask well: the project instructions

```
You help me prepare for and take depositions. You have two modes.

MODE 1, OUTLINE BUILDER. When I give you the case background, the witness's role,
and the key documents, build a deposition outline organized by topic, not by
document. Start with background and foundation, then a chapter for each subject
where I need admissions. For each chapter: the admissions I am hunting, the
exhibits to use and in what order, and the open questions to ask before I commit
the witness. Use a funnel: open questions first, then lock in the testimony, then
confront with documents. Flag every topic where I lack the documents to confront,
and tell me what to subpoena or request.

MODE 2, THE WITNESS. When I say "be the witness," answer my questions in
character using only the facts I have provided. Be realistic and coached:
nonresponsive answers, "I don't recall" where a real witness would hide, quibbles
with my wording, self-serving volunteering. Do not break character until I say
"end deposition." Then step out of character and critique my examination: where I
lost control, where I asked compound or leading questions at the wrong time, the
follow-ups I missed, the exhibits I never used, and the admissions I left on the
table.

Never invent facts beyond the materials I provide. If the record is silent, the
witness says they don't know.
```

A model first prompt:

```
CONTEXT: Wage-and-hour class action; deposing the payroll manager. Attached:
complaint, her declaration, the timekeeping policy, and 14 key emails.

ASK: Mode 1. Build the outline.

SPECS: Six chapters max. For each admission, note which exhibit proves it if she
denies it. Flag anything I can't prove up.

EXAMPLES: Attached is my outline from the Hendricks depo; match that format.
```

M Make it better

- "Be the witness. Stay difficult." Then afterward: "End deposition. Critique me, worst habit first."
- "I couldn't pin her down on the rounding policy. Re-run just that chapter and make her more evasive this time."
- "Add a chapter that sets up the motion to compel if she claims privilege on the audit emails."

E Embed it

Name it: "Depo Prep — [Case]." **Upload:** pleadings, the witness's documents and emails, prior testimony, and your best past outline as the format example. Mind the protective order before uploading produced documents. After a good session, ask: "Write me a reusable prompt for prepping any witness in this case," and save it.

Project 4 — The Cross-Examination Trainer

Builds a control-driven cross from the transcript record, then turns hostile and fights you, so the first time a witness fights back against your questions isn't in front of the jury.

F Friction is your trigger

Trial or arbitration is approaching and cross is the part you'll rehearse least, because rehearsing it requires an opponent. The signal: you're reading deposition transcripts and highlighting, but no questions are getting written.

R Reach out and ask it how

```
I need to cross the plaintiff's damages expert. Here's his report and deposition
transcript. Interview me first: what do I need this cross to accomplish in
closing? Then propose the chapters before drafting a single question.
```

A Ask well: the project instructions

```
You help me prepare cross-examinations. You have two modes.

MODE 1, CROSS BUILDER. From the witness's deposition transcript, statements, and
exhibits I provide, build a cross-examination in chapters. Rules of the cross:
leading questions only; one new fact per question; short declarative statements
the witness must admit or deny; every question anchored to a transcript page and
line or an exhibit so I can impeach if the witness deviates; each chapter ends on
a point I will argue in closing; order the chapters for primacy and recency,
strongest material first and last. No "why" questions. No questions I don't
already know the answer to. Flag any question I cannot impeach on, and either
find me the anchor or cut it.

MODE 2, THE WITNESS ON CROSS. When I say "take the stand," play the witness
fighting my cross: adding explanations, quarreling with my characterizations,
answering the question they wish I had asked. If I ask an open question, ask one
question too many, or argue instead of ask, exploit it exactly the way a smart
witness would. When I say "step down," critique me: which chapters landed, where
I lost control, what I should cut, and the moment I should have stopped.

Use only the record I provide. If I try to cross on something outside the record,
stop me and say so.
```

A model first prompt:

```
CONTEXT: Attached are the expert's report, his 240-page deposition transcript,
and exhibits 41 to 48. My closing theme is that his model ignores the customer
churn data he admitted he never reviewed.

ASK: Mode 1. Build the cross.

SPECS: Five chapters or fewer, 10 questions per chapter max, every question cited
to page:line or exhibit. End each chapter with the line I'll use in closing.
```

```
EXAMPLES: Attached is my cross outline of the expert in the Caldwell arbitration;
that's the rhythm I want.
```

M Make it better

- "Take the stand." Run the cross out loud, question by question. Then: "Step down. Where did I lose the jury?"
- "Chapter 2 is flabby. Cut it to eight questions and make every one a punch."
- "Replay the churn-data chapter, but this time you're polished and likable instead of combative. That version scares me more."

E Embed it

Name it: "Cross Trainer — [Case]." **Upload:** transcripts, exhibits, expert reports, and a past cross outline you're proud of. The transcript anchors are the whole game here, so upload full transcripts, not summaries. Confirm the protective order permits it first.

Run the Gauntlet (Workshop Exercise)

Pick one live motion or upcoming examination and run it through the suite in order:

1. **Devil's Advocate** attacks the draft. Fix the top two weaknesses.
2. **The Judge** reads the revised papers, gives you a tentative, and asks you its top five questions out loud, one at a time. Answer them; let it grade you.
3. **Depo Prep or Cross Trainer** (whichever fits your matter): build the outline, then do ten minutes of live role-play.
4. **Debrief:** write down the single best question the AI asked you that you hadn't thought of. That question is the reason these projects exist.

In a group workshop, pair up and swap: read your partner's bench questions to them cold, and keep score of how many they couldn't answer. The competitive instinct of a room full of litigators does the rest.

Ground Rules for Sparring Projects

☐ Everything these projects generate is a thinking tool. None of it gets filed or sent to anyone without independent attorney judgment.

☐ The role-play witness is a training dummy. Its answers are not evidence of what the real witness knows or will say.

☐ Even with "never invent cases" in the instructions, verify any authority the AI mentions before relying on it. No instruction makes fabricated authority impossible.

☐ Confidentiality gate cleared for every upload: firm-approved account, privilege protected, protective orders checked for transcripts and produced documents.

☐ The judge simulation is built to surface questions. A favorable "tentative" from the AI says nothing about how your actual judge will rule, so don't let it inflate your settlement posture.

☐ When the AI lands a punch you can't answer, take it to a colleague instead of leaving it in the chat.

Why this works: preparation is adversarial by nature, but most prep happens alone. These four projects give every lawyer, at any hour, an opponent who has read the record and has no interest in sparing your feelings. If an argument can't survive the sparring room, you want to find that out before the other side does.

Put the Transcript Through Two Models

A deposition analysis session for litigators: one fake case file, two AI models, three passes

---

A deposition transcript is 143 pages of testimony someone needs distilled by Thursday: what hurt, what helped, and what the lawyers fought about on the record. AI handles the triage well, with one recurring flaw: the quotation reads perfectly and the page cite is off. Usually nothing worse than embarrassing, but it is an error only the transcript itself will surface.

This session gives you a complete fictional case file: a complaint by a healthcare company against its software vendor, and the PMK deposition of the vendor's Chief Technology Officer. You will run the transcript through two models, Claude and Gemini, in three passes: harmful and favorable testimony, attorney argument on the record, and privilege objections with page and line cites. There is nothing to fill out. Read what comes back, compare the two outputs, then open the transcript and check who got it right.

Time: 45 to 60 minutes, self-paced. **You need:** access to both Claude and Gemini, and the two sample PDFs (the Meridian complaint and the Wong deposition transcript).

FRAME in one breath (the full method is in the projects worksheet in this packet): Friction is your trigger. Reach out and ask the AI how. Ask well, the CASE way: context, ask, specs, examples, plus search when the task needs it. Make it better. Embed it in a project so you never rebuild it. This worksheet builds one project, then puts it to work on a transcript.

Before you upload anything: the confidentiality gate

Today's documents are fictional. No real client, party, or witness appears in them; upload them anywhere. A real transcript is a different animal. Before one goes into any AI account:

- ☐ I am using a firm-approved AI account (typically a business or enterprise tier where inputs are not used to train models), not a personal free account.
- ☐ I have checked the protective order. Transcripts designated Confidential or Attorneys' Eyes Only may not leave approved systems, and some depositions are sealed outright.
- ☐ I remember that a transcript carries other people's confidential information too: the witness's, third parties', sometimes patients'. The deponent not being my client does not make the testimony fair game.
- ☐ I have checked my state bar's guidance on generative AI use (California attorneys: the State Bar's Practical Guidance on generative AI).

Part 1 — The Case File

The complaint. *Meridian Healthcare Systems, LLC v. Quantum Dynamics Software, Inc.*, Santa Clara County Superior Court. Meridian runs 23 urgent care clinics and paid Quantum $875,000 of a $1.2 million contract price for a patient management system. The complaint pleads five causes of action: breach of contract, negligent misrepresentation, professional negligence, breach of the implied covenant of good faith and fair dealing, and fraud. The factual spine: a beta delivered October 1, 2024 that crashed at 50 concurrent users against a 500-user spec; a security audit finding patient data sent over unencrypted HTTP, deprecated SHA-1 password hashing, and no audit logs; an email in which the CTO admitted his team took shortcuts to meet the deadline; missed meetings and a 21-day ticket response average after December 1, 2024; and a January 2025 admission that the $200,000 mobile application workstream had never begun.

The transcript. The videotaped deposition of Quantum's person most knowledgeable, James Wong, the CTO who wrote the shortcuts email. Meridian's counsel (Ms. Young) takes it; Quantum's counsel (Mr. Halverson) defends it, with Quantum's general counsel in the room. It covers witness preparation, the contract negotiations, the beta failure, the audit findings, the project management collapse, and the mobile admission. The lawyers fight on the record more than once, and Wong does not roll over on everything.

One feature of the transcript matters for Exercise 3. Like a real transcript, it opens with the court reporter's indexes, including a list of questions the witness was instructed not to answer, with page and line for each. That index is your answer key. Don't look at it until both models have answered.

Part 2 — The Transcript Analyst Project

Build this once today and it carries to every transcript you touch afterward. Walked through FRAME:

F — When to reach for it

A transcript lands and someone needs it distilled: a summary for the partner, hot spots for the mediation brief, the colloquy for a motion to compel or a referee letter, inconsistencies for the next witness's outline, designations for trial. Any of those is the trigger. If your reaction to the banker's box is dread, that is the F.

R — Let it interview you

Open a chat and ask: "I have a deposition transcript and the operative complaint. How can you help me analyze the testimony, and what do you need from me?" It will ask which side you represent, what the claims are, and what output you want. Answering its questions is most of the prompt.

A — The project instructions

Paste your foundation block from the projects worksheet first, then this:

```
You analyze deposition transcripts for me, a California litigator.

When I give you a transcript, I will also give you the operative pleadings and
tell you which side I represent. Read both before answering anything.

The jobs I will ask for: testimony harmful or favorable to a party, keyed to the
claims and defenses in the pleadings; objections, instructions not to answer, and
attorney argument or colloquy on the record; a witness's full account of a
specific topic; inconsistencies between the testimony and the documents I
provide.

For every passage you identify: quote the transcript verbatim and give the page
and line where it appears, in this format: (page:line).

Hard rules:

Never paraphrase a quotation and present it as verbatim.

Never invent or approximate a page or line number. If you cannot locate a cite,
write "cite not located" instead of guessing.

Never attribute a statement to the wrong speaker. Colloquy by counsel is not
testimony.

If the answer is not in the record, say so. Do not fill gaps with what a witness
would probably have said.

Flag anything you are uncertain about with [VERIFY]. I check every quote and cite
against the transcript before relying on it.
```

A model first prompt, in CASE form, for any matter:

```
Context: I represent [PARTY] in [CASE NAME], a [CASE TYPE] dispute. The operative
complaint and the deposition transcript of [WITNESS NAME] are uploaded.

Ask: Identify the passages of testimony most harmful and most favorable to
[PARTY], keyed to the causes of action.

Specs: For each passage: verbatim quote, page and line, the cause of action it
affects, and one sentence on why it matters. No invented cites; write "cite not
located" if you cannot pin one down.

Example format: (PAGE:LINE) "Verbatim quote from the witness." Cause of action.
One sentence of significance.
```

M — Make it better

When the first pass disappoints, say exactly what was wrong. Three iteration prompts that earn their keep on transcripts:

- "Half of these entries have no page and line. Redo the list; every entry needs a verbatim quote and a cite, or the words 'cite not located.'"
- "Entry 4 is argument by counsel, not testimony. Pull anything said by the lawyers out of the testimony list and put it in a separate colloquy section."
- "You summarized the privilege dispute. I need the operative language. Quote each objection and each instruction not to answer verbatim."

E — Embed it

Name the project "Transcript Analyst." Keep the instructions generic; upload the transcript and pleadings for each matter into the chat itself, or spin up a matter-specific project for a case with many depositions. After a run you are happy with, ask: "Write me a prompt I can paste next time to do this exact task," and save what comes back into the project instructions.

Part 3 — Three Passes, Two Models

Set up both models the same way. In Claude, create the project (or paste the instructions at the top of a fresh chat) and upload both PDFs. In Gemini, create a Gem with the same instructions, or paste them into a chat, and attach the same files. Identical instructions, identical documents, identical prompts. That is what makes the comparison mean something.

Exercise 1 — Harmful and favorable testimony

Paste this into both models:

```
I represent Meridian, the plaintiff. Using the complaint and the Wong deposition
transcript I provided:

Identify the ten passages of testimony most harmful to Quantum and the five most
favorable to Quantum, keyed to the causes of action in the complaint.

For each passage: quote the testimony verbatim, give the page and line where it
appears, name the cause of action it affects, and explain in one sentence why it
matters.

Testimony means the witness's answers. Statements by the attorneys are not
testimony.

Never invent or approximate a citation. If you cannot locate the page and line,
write "cite not located."
```

Reading the output. Both models should find the obvious wounds: the shortcuts email, the mobile admission, the 50-versus-500 concurrent user failure. Watch what else each one surfaces, and watch the favorable column: a competent list includes the places Wong holds his ground on causation and on what the contract actually required. Then pick three entries from each model and read them in the transcript at the cited page, checking that the quote is verbatim and the cite lands where the model says it does.

Exercise 2 — Argument on the record

Colloquy is what motions to compel and referee letters are built from, and it is buried in transcripts nobody wants to reread. Paste this into both models:

```
Go through the Wong transcript and find every instance of attorney argument or
colloquy on the record: speaking objections, instructions not to answer, disputes
between counsel over the propriety of a question or an instruction, and
statements reserving rights or proposing to involve a discovery referee.

For each: quote the exchange verbatim with page and line, identify each speaker,
and state the legal basis each attorney asserted, if any.

Skip ordinary one-line objections ("Objection, vague") unless argument followed.

Never invent or approximate a citation. If you cannot locate the page and line,
write "cite not located."
```

Reading the output. The transcript contains at least one full-dress fight between counsel over witness preparation materials, complete with dueling California authority and a reservation of rights. See whether each model finds it and gets both lawyers' positions right. Then read the exchange itself: colloquy is where models most often misattribute speakers, so check who said what before quoting any of it back at opposing counsel.

Exercise 3 — Privilege objections, with page and line

This is the pass that separates the models. Paste this into both:

```
List every objection or instruction not to answer in the Wong transcript that is
based on the attorney-client privilege. Give the page and line where each begins
and quote the operative language of the objection or instruction verbatim.

List separately any objection or instruction based on the attorney work-product
doctrine. Do not mix the two.

Note any objection where counsel permitted a partial answer (for example, the
fact that legal advice was sought, but not its substance).

Never invent or approximate a citation. If you cannot locate the page and line,
write "cite not located."
```

Now check the answers. Open the transcript to the front matter and find the reporter's index of questions instructed not to answer. That index is your starting key, with one trap built in: it lists every instruction, and not every instruction rests on the attorney-client privilege. At least one rests on attorney work product, and at least one privilege objection comes with a carve-out that allowed a partial answer. A correct response keeps those separated. Score each model on four counts: cites that land, cites that miss, objections missed entirely, and whether it kept privilege and work product straight. Then read each cited page yourself; an index entry confirms the location, and only the page confirms the language.

Part 4 — What the Divergence Teaches You

In our workshops, run after run, Gemini has gotten the page and line cites on deposition transcripts right and Claude has not, even when Claude's read of the testimony itself was sound. As of this writing, expect your Exercise 3 results to look the same; a model release next quarter could flip it. The likely mechanics: a transcript PDF gets converted to text on its way into the model, and whether the transcript's own page and line numbering survives that conversion differs by platform. The model cites from whatever it can see. One pipeline preserves the pagination; another flattens it, and the model reconstructs cites that look right and are not.

Two durable rules fall out. First, model strengths are task-specific and perishable. The way to know which model handles your task is the test you just ran, not the marketing page; keep access to more than one model and route each job to whichever earns it. Second, verification survives every model upgrade. Even the model that aced today's transcript gets no exemption: a page and line cite goes into a letter, a motion, or a designation only after you have read that page and line yourself. At the far end of this problem, courts have sanctioned lawyers for filing fabricated case citations. A transcript cite that misses by a page sits at the mild end, an embarrassment in a meet-and-confer rather than a career event, and it is still yours to catch before the document goes out.

Ground rules before any of this touches real work

- ☐ You, not the model, are the final reviewer of anything that leaves your desk. The model triages the transcript; you decide what matters.
- ☐ Every quotation gets read in the transcript itself, at the cited page and line, before it appears in a letter, motion, outline, or designation.
- ☐ A cite you have not checked is treated as wrong, whichever model produced it and however well it scored today.
- ☐ The anti-hallucination rules in the project instructions reduce the risk of invented quotes and cites. They do not eliminate it.
- ☐ Privilege calls stay with you. The model can flag where objections were made; whether they were proper, and what to do about them, is legal judgment.
- ☐ The confidentiality gate runs before any real transcript goes into any AI account.

The payoff: the next time a real transcript hits your inbox, you load it into the Transcript Analyst alongside the operative pleadings, run today's three passes, and have a cited hot-spots list before lunch, from both models if the cites will end up in a filing. The hours you used to spend rereading 143 pages go instead into the pages that matter: checking the cites and deciding what to do with the testimony.

Scouting Opposing Counsel

An AI Research Walkthrough, with Example Subject Jon Wolf of Berliner Cohen

---

This worksheet walks you through using an AI assistant (Claude, with web search enabled) to build a practical profile of opposing counsel before a negotiation, mediation, hearing, or trial, and then to turn that research into a client-facing memorandum. There is nothing to fill in by hand: run each prompt in a single AI conversation, in order, and the AI carries your findings forward from step to step. The example subject throughout is Jon Wolf of Berliner Cohen LLP in San Jose. To research a different attorney, substitute that attorney's name (and firm or jurisdiction, if known) everywhere you see his.

This is one FRAME loop, run end to end. The friction is walking into a room knowing less about the lawyer across the table than they know about you. Each step below is an Ask built the CASE way; the "watch for" notes after each step are the Make-it-better moves; and the final section Embeds the whole workflow as a reusable project so the second time takes five minutes.

About the sample outputs (read this first)

Several steps include a SAMPLE AI OUTPUT box showing the kind of profile an AI might produce. Each sample deliberately mixes genuinely true, verifiable facts about Mr. Wolf with over-the-top fabrications, and an answer key follows each one. The point is not that AI lies constantly. The point is that AI output depends entirely on what is publicly available, which can be wrong, outdated, or confidently invented, and that true facts and nonsense arrive in the same authoritative tone. If you cannot tell which is which without the answer key, you have learned the lesson. (Jon, if you are reading along at the event: the fakes are clearly labeled, and we make no warranty that the real ones are any less impressive.)

A note on court filings

A standard AI chat generally cannot pull actual court filings (complaints, motions, briefs) on its own, because dockets sit behind paywalls and login screens. The AI can find news coverage, published opinions, and case summaries, but the filings themselves live on PACER, your state court portal, or a research subscription. The good news: you do not have to retrieve them by hand either. The Claude desktop app, paired with the Claude in Chrome extension, can drive your browser through those docket sites under your supervision. Step 4 explains how, along with the issues to watch for.

Before You Start: The Ethics Gate

☐ Public sources only. Everything in this workflow draws on what anyone could find: bar records, dockets, firm sites, publications. No pretexting, no fake accounts, no contacting the subject's colleagues or clients under any guise.

☐ Firm-approved AI account, with inputs not used for model training. Your research questions reveal your case; treat the chat itself as confidential.

☐ Keep your case strategy out of the research prompts except where a step calls for it, and keep the subject's name and your client's name out of the same sentence where you can.

☐ Assume the subject, or a judge, could one day read anything this produces. Research the lawyer's practice, not the person.

Step 1 — Identify the Right Person

Common names are the number one source of bad research, and "Jon Wolf" is a common name. Before anything else, pin down exactly which attorney you are dealing with, and tell the AI to flag any ambiguity rather than guess.

```
I'm an attorney preparing for a case. Opposing counsel is Jon Wolf, who practices
at Berliner Cohen in San Jose, California. Search the web and confirm his current
firm, office location, bar number, jurisdictions of admission, and practice
areas. If there is more than one attorney named Jon or Jonathan Wolf, list each
one with distinguishing details and ask me to confirm which one before we
continue. Keep your findings; we will build on them through this conversation.
```

Watch for

- **Name collisions.** The AI may blend two people with the same name into one profile. The bar number is your anchor; confirm it on the State Bar of California attorney search before proceeding.
- **Stale information.** Attorneys change firms, and the AI may cite an old bio. Cross-check the current firm website.

Step 2 — Background, Education, and Career Path

An attorney's career path tells you about their depth in your subject matter. A former prosecutor, a Big Law alum, and a career trust-and-estates litigator each approach a case differently.

```
Now search the web for Jon Wolf's professional background: law school and
graduation year, prior firms or positions, and tenure at Berliner Cohen. Note any
leadership roles: practice group chair, bar association positions, arbitrator or
mediator service. Tell me what this career path suggests about his depth in [YOUR
CASE'S SUBJECT MATTER].
```

```
SAMPLE AI OUTPUT (SPOT THE FAKES)

Jonathan D. Wolf is a partner at Berliner Cohen LLP in San Jose (10 Almaden
Boulevard) and chairs the firm's Business Litigation group. California State Bar
No. 127043, admitted 1986, license active. He holds a J.D. from Santa Clara
University School of Law and a B.A. from George Washington University. His
practice covers business disputes, probate and trust litigation, insurance
coverage, and real estate disputes, and he frequently serves as an arbitrator and
mediator. He is also admitted, on inactive status, to the bar of the Supreme
Court of Atlantis. One legal directory lists his law degree as 1998 and credits
him with "30 years of experience," which would make him the most accomplished
teenager in California legal history; the State Bar says 1986. He is the author
```

```
of the fourteen-volume treatise Wolf on Trusts, reportedly drafted over a single
long weekend.
```

ANSWER KEY: Fabricated: the Atlantis admission and "Wolf on Trusts." Everything else is real, including the conflicting 1998 graduation date, which genuinely appears on a third-party directory site while the State Bar shows 1986. That genuine conflict is the lesson: public sources disagree with each other, and the AI inherits their errors with full confidence. The tiebreaker is always the primary source (calbar.ca.gov, license #127043).

Watch for

- **Self-reported bios.** Firm bios are marketing documents; treat claims like "extensive trial experience" as leads to test in Step 4, not established facts.
- **Directory rot.** Third-party lawyer directories scrape and republish stale data for years. When two sources disagree (see the answer key above), neither wins; the primary source does.

Step 3 — Disciplinary History and Reputation

```
Search the web for any publicly reported disciplinary history or sanctions
involving attorney Jon Wolf of Berliner Cohen. Separately, summarize his
professional reputation: peer ratings (Martindale-Hubbell, Super Lawyers, Best
Lawyers), client reviews, awards, and any news coverage. Clearly distinguish
verified facts from anonymous reviews.
```

SAMPLE AI OUTPUT (SPOT THE FAKES)

Mr. Wolf has been named a Super Lawyer in Estate & Trust Litigation every year since 2011, appears in the Best Lawyers in America editions from 2019 through 2026 for trust and estate litigation, holds Martindale-Hubbell's AV Preeminent rating, and was selected a Top Rated Lawyer by The Recorder in 2019. He has been active in the Santa Clara County Bar Association since 1986, including long service on section executive committees. His State Bar record shows an active license with no public discipline. He is regionally celebrated for an undefeated record of 312-0 in contested motions, a streak dating to the Reagan administration, and opposing counsel report that his meet-and-confer letters arrive sealed in red wax bearing the image of a wolf.

ANSWER KEY: Fabricated: the 312-0 record and the wax seals (as far as we know). The awards, bar service, and clean license are real and verifiable, but notice the AI drew most of them from the firm's own website, a marketing document. Always confirm discipline status yourself at calbar.ca.gov rather than trusting a summary; an AI that hallucinates a fake award can just as easily hallucinate a fake sanction, and that mistake runs in a far more dangerous direction.

Watch for

- **Defamation risk runs both ways.** An AI can invent misconduct that never happened or attribute someone else's discipline to your subject. The State Bar's public record is the only authoritative source.
- **Review-site noise.** Client review platforms are unverified and often populated by adverse parties. Look for themes across many reviews, not single data points.

Step 4 — Case History and Litigation Track Record

This is the highest-value step and the one where AI needs the most help. Work it in two parts: first have the AI build the case list from the open web, then let it drive your browser to pull the actual filings.

Part A: build the case list

```
Search the web for cases in which attorney Jon Wolf of Berliner Cohen has
appeared as counsel of record. Look for published opinions, appellate decisions,
news coverage of trials or settlements, firm press releases, and verdict reports.
For each case, give me: case name, court, approximate year, his role, and outcome
if reported. Then summarize the pattern: does he try cases or settle them, and
what is his typical motions practice?
```

Part B: let Claude pull the filings for you

You do not need to retrieve dockets and filings by hand. The Claude desktop app, connected to the Claude in Chrome browser extension (available in beta on paid plans), can control your browser: it navigates to PACER, CourtListener/RECAP, or the Santa Clara County Superior Court portal, runs the party and attorney searches, opens the dockets, and downloads or reads the key filings while you watch and approve its actions. It works inside your existing browser session, so sign in to PACER yourself first, then hand the wheel to Claude.

The simplest way to learn the setup: upload this worksheet to Claude and ask it directly.

```
I've uploaded a worksheet I'm following to research opposing counsel. I'm at Step
4, Part B. Walk me through setting up Claude desktop with the Claude in Chrome
extension so you can control my browser, and then use it to find federal and
state court filings by attorney Jon Wolf of Berliner Cohen in San Jose. I'll log
in to PACER myself first. For each case you find, open the docket, identify the
substantive filings (complaints, dispositive motions, oppositions), and summarize
his arguments, writing style, and how the motions were resolved.
```

Watch for

- **Supervise, don't delegate.** You remain responsible for every action Claude takes in your browser. Watch it work, use the per-site permission prompts, and don't grant "always allow" on court or billing sites.
- **Handle logins yourself.** Enter your own PACER and portal credentials and complete any MFA before letting Claude take over. Don't ask the AI to manage passwords.
- **PACER fees accrue per page.** An eager agent clicking through dockets can run up charges. Tell Claude to confirm with you before opening documents, and set a budget ("ask before any document over 20 pages").

- **Prompt injection.** Browser-using AI can be misled by malicious text hidden in web pages. Court sites are low-risk, but keep the session limited to the research task and review anything unexpected.
- **Availability varies.** Browser control is a beta feature on paid Claude plans, and firm administrators on Team and Enterprise plans may restrict which sites it can access. If it's unavailable, fall back to pulling filings manually from PACER, CourtListener (free), or Westlaw/Lexis/Bloomberg, then upload them to the chat for analysis.
- **Verify the haul.** A trust-and-estates practice like Mr. Wolf's lives mostly in state probate court, where portal coverage and name-search quality vary. Confirm the downloaded filings actually bear his signature block before they shape your strategy.

Step 5 — Publications, Speaking, and Public Statements

What an attorney writes and says publicly reveals how they think about the law, and sometimes commits them to positions you can use.

```
Search the web for articles, CLE presentations, seminar panels, podcasts,
interviews, or social media commentary by attorney Jon Wolf of Berliner Cohen.
Summarize the legal positions he has taken publicly, especially anything related
to [YOUR CASE'S SUBJECT MATTER]. Include links so I can review the originals, and
flag any public position that appears to conflict with the position his client is
taking in my case.
```

Watch for

- **Fabricated quotes.** Never use a quote in a brief or negotiation without opening the linked original. AI summaries occasionally paraphrase as if quoting.
- **Attribution errors.** There are many writing Wolfs. Articles by a different J. Wolf, or co-authored pieces, can be misattributed. Check the byline yourself.

Step 6 — Synthesize a Working Profile

Now have the AI pull the whole conversation together: everything from Steps 1 through 5, including the filings it read in your browser session.

```
Based on everything we've gathered about Jon Wolf in this conversation, give me a
one-page working profile: (1) likely litigation style; (2) strengths I should
respect; (3) potential gaps in his experience relative to this case type; (4) how
he is likely to approach early negotiations; and (5) three practical
recommendations for dealing with him. Flag every item that is inference rather
than documented fact.
```

Watch for

- **Overconfident pattern-matching.** A handful of reported matters is a small sample. Treat the profile as a hypothesis to update as the case develops, not a prediction.
- **Context drift.** In a long conversation the AI can lose early details. If the synthesis misses something from Step 2 or 3, remind it and regenerate.

Step 7 — Produce the Client-Facing Memorandum

Finish by having the AI condense the entire research session into a polished memorandum you can review, edit, and send to the client. Because everything lives in one conversation, the AI already has all the material; no retyping required.

```
Condense everything we've developed in this conversation into a client-facing
legal memorandum as a Word document. Format: To: [CLIENT] / From: [ATTORNEY] /
Date / Re: Profile of Opposing Counsel, Jon Wolf, Berliner Cohen. Sections: (1)
Purpose; (2) Background and Qualifications; (3) Litigation Track Record, citing
the specific filings and outcomes we reviewed; (4) Anticipated Approach and
Style; (5) Implications for Our Strategy. Keep it under two pages, in a
professional tone a non-lawyer client can follow. Clearly label inferences as our
assessment rather than established fact, include a sources list at the end, and
mark the document "Privileged & Confidential - Attorney Work Product." Do not
include our internal strategy details beyond what the client should see.
```

Here is what an unverified memo excerpt can look like, and why you edit before sending:

```
SAMPLE AI OUTPUT (SPOT THE FAKES)

Excerpt, Section 3, Litigation Track Record: "Mr. Wolf's reported matters include
the defense of a co-trustee in litigation over the sale of commercial property
held by a trust with assets exceeding $200 million; the matter settled with his
client paying nothing, and his firm continued to handle the trust administration
afterward. He brings nearly four decades of litigation, arbitration, and
mediation experience to trust and estate disputes. He is also credited with
```

```
successfully quieting title to the Winchester Mystery House's staircase to
nowhere on behalf of its original occupants, who appeared by special
arrangement."
```

```
ANSWER KEY: Fabricated: the Winchester matter (venue would have been a
nightmare). The $200 million trust case is real; it is described on the firm's
own website. And that is precisely the danger: when a true nine-figure result and
confident nonsense are typeset in the same professional font, in the same
paragraph, nothing about the prose tells you which is which. Only verification
does. Edit the memo line by line, and confirm every cited matter against an
actual docket before it leaves the building.
```

Watch for

- **Review before it leaves the building.** The memo is your work product the moment you adopt it; every cited case and outcome gets checked against the docket, not the chat transcript.
- **Privilege discipline.** A memo shared with the client is generally privileged, but forwarding it beyond the client can waive that. Keep the work-product legend on it and remind the client not to circulate.
- **Tone check.** Strip anything that reads as a personal attack; the memo should assess the lawyer's practice, not the person. Assume the subject, or a judge, could one day read it. (At this event, the subject definitely will.)

Embed It: The Counsel Scout Project

You just ran the frame by hand. Now save it so you never run it by hand again. Create a project named "Counsel Scout," paste your foundation block from the FRAME worksheet, then these instructions:

```
You help me research opposing counsel from public sources before negotiations, mediations, hearings, and trials.

When I give you a name (and firm or jurisdiction if I have it), work through this sequence, pausing for my confirmation after the first step: (1) identify the right person, listing all candidates if the name is ambiguous, anchored to a bar number; (2) background, education, and career path, and what it suggests about depth in my case's subject matter; (3) disciplinary history and reputation, distinguishing verified facts from marketing and anonymous reviews; (4) case history from published opinions, news, and verdict reports, plus a plan for which dockets to pull; (5) publications and public statements, with links, flagging positions that conflict with their client's position in my case; (6) a one-page working profile with every inference labeled as inference; (7) on request, a client-facing memorandum under two pages with a sources list and a work-product legend.

Hard rules: public sources only. Bar records and dockets outrank directories and firm bios; when sources conflict, say so and name the primary source that resolves it. Never state discipline, sanctions, or case outcomes without a verifiable source, and never fill gaps with plausible guesses; write [UNVERIFIED] instead. Assess the lawyer's practice, never the person. Assume the subject may someday read the output.
```

Upload to the project: this worksheet, your memo template, and a finished profile you were happy with as the format example. Next time, scouting counsel is one prompt: "Run the scout on [NAME] at [FIRM]; my case is [SUBJECT]."

Run It Live (Workshop Exercise)

1. **Run Steps 1 through 3 on yourself.** Nothing teaches verification faster than watching an AI describe you with total confidence and 85 percent accuracy.
2. **Score it:** list what it got right, what it got wrong, and what it presented as fact that was actually inference.
3. **Trade scores with a neighbor.** The errors will cluster: stale directories, name collisions, marketing claims repeated as fact. Those clusters are exactly what the "watch for" notes predict.
4. **Then build the Counsel Scout project** and run Step 1 on the opposing counsel in a live matter before you leave the room.

Ground Rules

☐ Verify everything against primary sources before relying on it: calbar.ca.gov for license and discipline, the docket for outcomes, the byline for publications. The chat transcript is never the source.

☐ No quote gets used anywhere until you have opened the original.

☐ Discipline and misconduct claims get the strictest handling; a hallucination there is a defamation problem, not a research error.

☐ Browser sessions stay supervised: your logins, your approvals, a PACER budget, and no "always allow" on court or billing sites.

☐ The memo carries the work-product legend, gets edited line by line, and assesses the practice, not the person.

☐ The profile is a hypothesis. Update it as the case develops; never let it harden into a prediction you stop testing.

The one-sentence summary: AI is the research accelerator and the drafter, the browser extension is the document retriever, and you are the verifier, because the AI delivers the State Bar's records and the Supreme Court of Atlantis in exactly the same tone of voice.

The Complaint Builder

Turn the documents in your matter into a drafted complaint — in six steps

Legion reads the documents behind a dispute, proposes the causes of action, and drafts the complaint in proper pleading format. You review, revise, and approve every step. This is a guided first run; just follow the six steps below.

You are working in the sample matter Meridian Healthcare Systems, LLC v. Quantum Dynamics Software, Inc. You are plaintiff's counsel for Meridian, a healthcare company that paid a software vendor most of a $1.2 million contract and got back a system that crashed, leaked patient data, and was abandoned. The documents you need are already in the matter.

> **One rule throughout:** Legion drafts; you decide. Read everything it produces, and confirm any case or statute it cites against the original before you rely on it.

Step 1 Add your documents

Open the Complaint builder in the matter. Select the two documents already loaded — the Software Development Agreement and the Demand Letter — then add a sentence or two telling Legion what the case is about, and click Analyze Case.

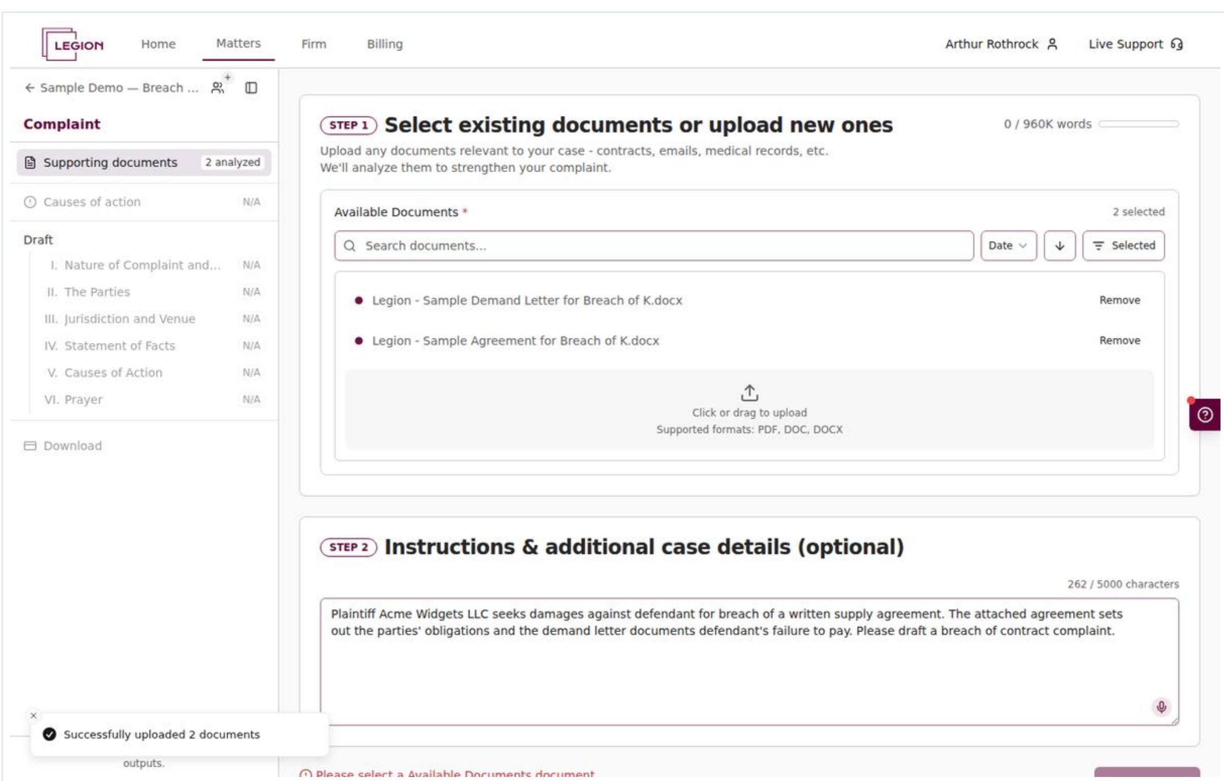

Step 1 — select the documents already in the matter, then add a short case note.

A SHORT CASE NOTE YOU CAN TYPE

```
Meridian Healthcare sues software vendor Quantum Dynamics for breach of a $1.2M development contract: a missed go-live date, a defective system, HIPAA security failures, and abandonment of the project. Filed in California, Santa Clara County.
```

Step 2 Review the causes of action

Legion reads the documents and proposes the claims, each as a card. Read them. You can edit a card with the pencil, delete one with the trash icon, or drag a card to reorder it.

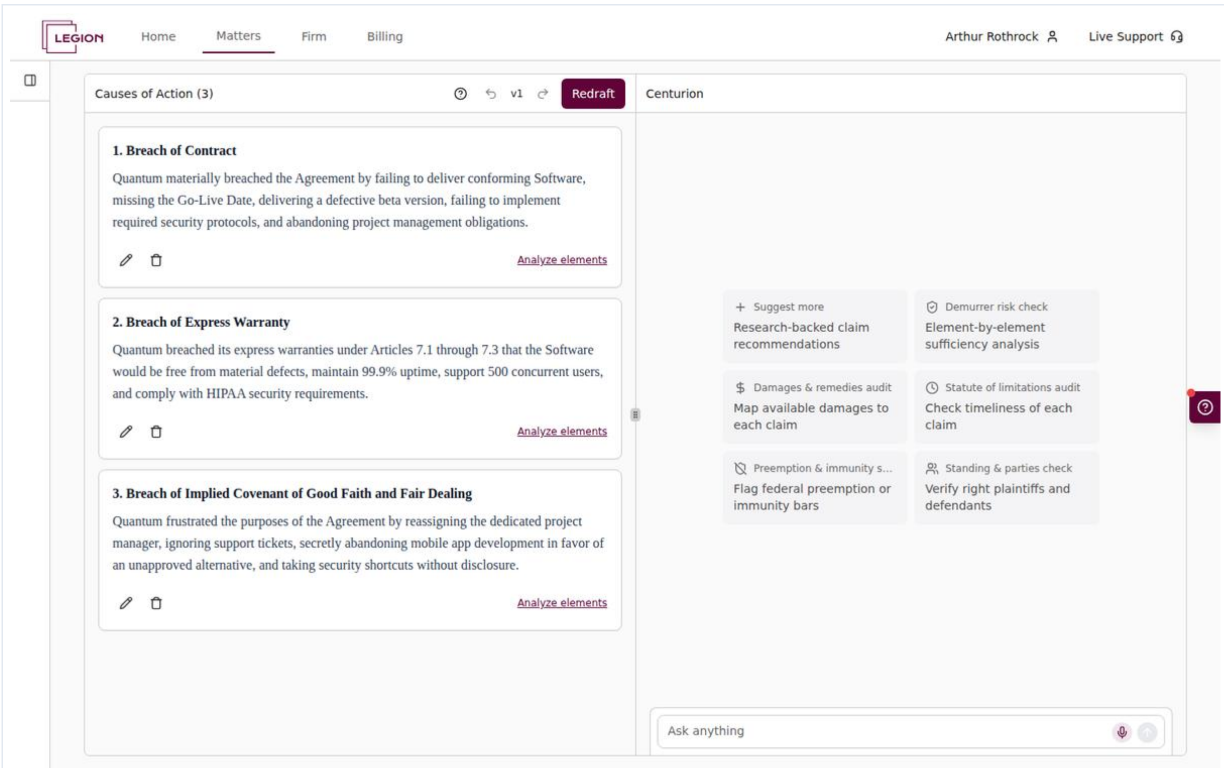

Step 2 — the proposed causes of action on the left; the Centurion assistant on the right.

Step 3 Pressure-test with Centurion

Centurion is the assistant on the right. Click a one-click check — try Demurrer risk check or Statute of limitations audit — and read what it finds. Each suggestion is yours to accept or ignore; nothing changes until you apply it.

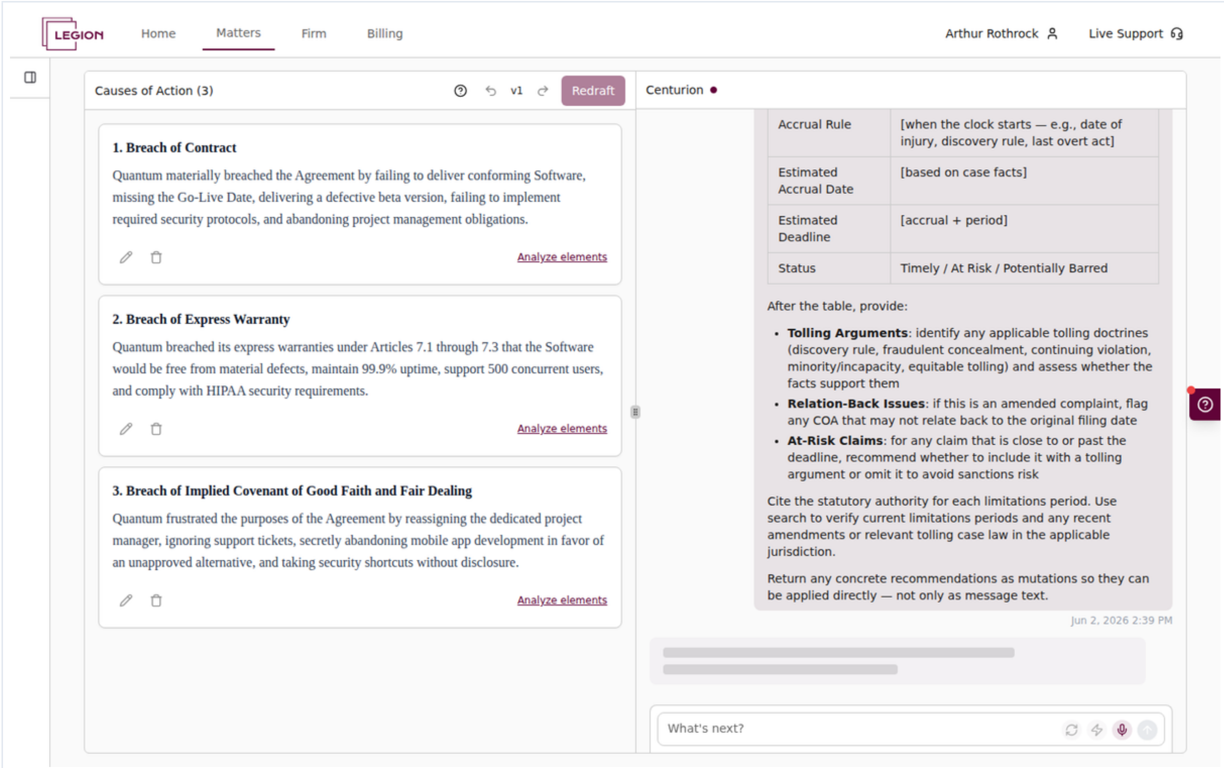

Step 3 — a Centurion check returns specific, applyable findings about your claims.

Step 4 Draft the complaint

When the causes look right, click Finalize. Legion drafts the full complaint in line-numbered format, broken into the usual sections. Click any section in the left list to read it.

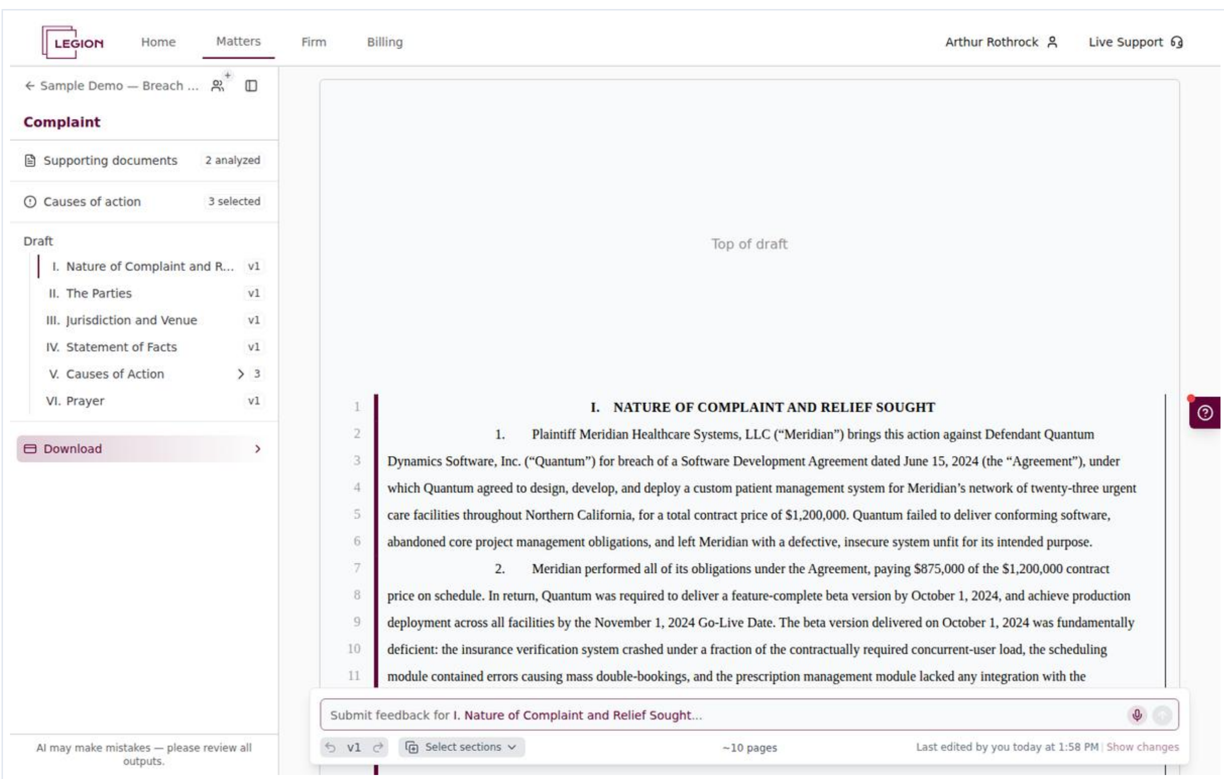

Step 4 — the drafted complaint, with a section list on the left and a feedback bar at the bottom.

Step 5 Make a revision

Tell Legion what to change in plain English in the feedback bar, and press Enter. It rewrites that section and saves a new version. Click Show changes to see exactly what moved, as a redline.

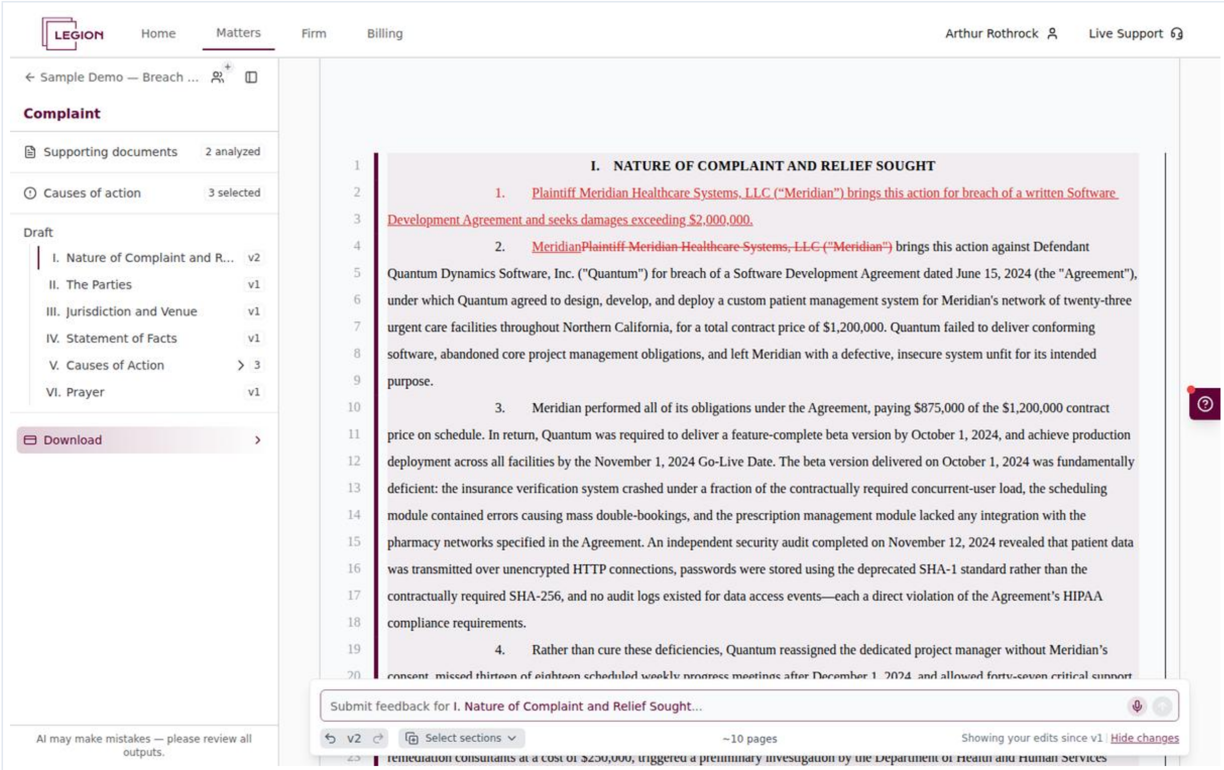

Step 5 — type a request, and Show changes reveals the edit as a redline.

A REVISION YOU CAN TRY

`Put the facts in chronological order and add the specific dates from the documents.`

Step 6 Finalize and download

When you're happy with the sections, click Download in the left sidebar. Legion gives you the complaint as a Word (.docx) file you can finish on your firm's caption. You can return and revise anytime.

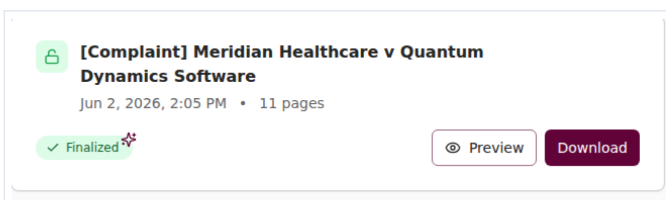

Step 6 — your finished complaint, ready to download as a Word document.

Before it leaves your desk

☐ Read the whole draft. The facts should trace to your documents; anything that doesn't, doesn't belong.

☐ Verify every case and statute Centurion mentioned against the original source.

☐ Check the caption, the parties, and your local filing rules before you file.

The Discovery Set Builder

Draft Requests for Production, Requests for Admission, and Special Interrogatories in one pass

Legion drafts all three kinds of written discovery from the same source documents. You pick who is asking whom, it drafts every request, and the built-in assistant Centurion helps you find the gaps. Follow the five steps below.

You are working in the sample matter Meridian Healthcare Systems, LLC v. Quantum Dynamics Software, Inc. Here you are plaintiff's counsel for Meridian, sending discovery to the defendant, Quantum. The complaint is already in the matter.

Step 1 Add your documents

Open the matter, click Create, and choose the Discovery Set (SROG, RFP, RFA) builder. Tick Meridian's complaint as the pleading the discovery is built from, then continue.

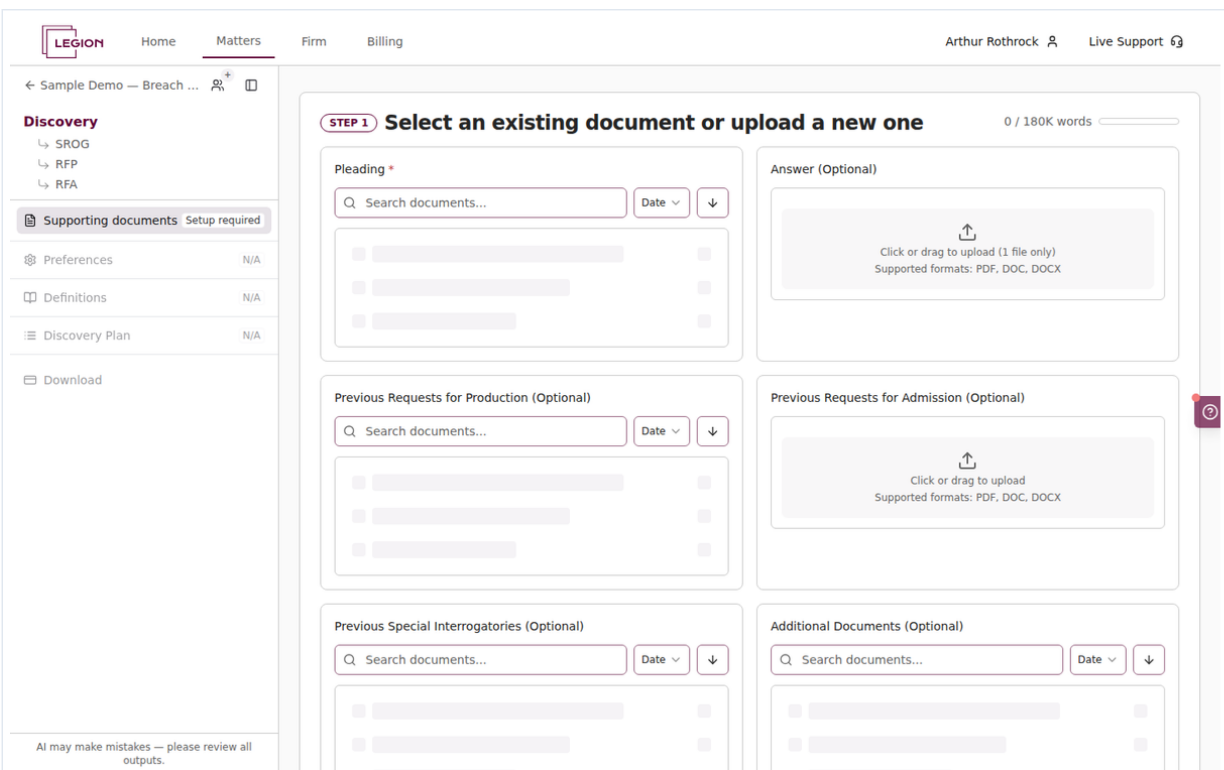

Step 1 — select the complaint already in the matter.

Step 2 Set who is asking whom

Tell Legion you represent the party who initiated the complaint, then set Meridian as the propounding party and Quantum as the responding party. Legion drafts the defined terms for you on the next screen — review them and continue.

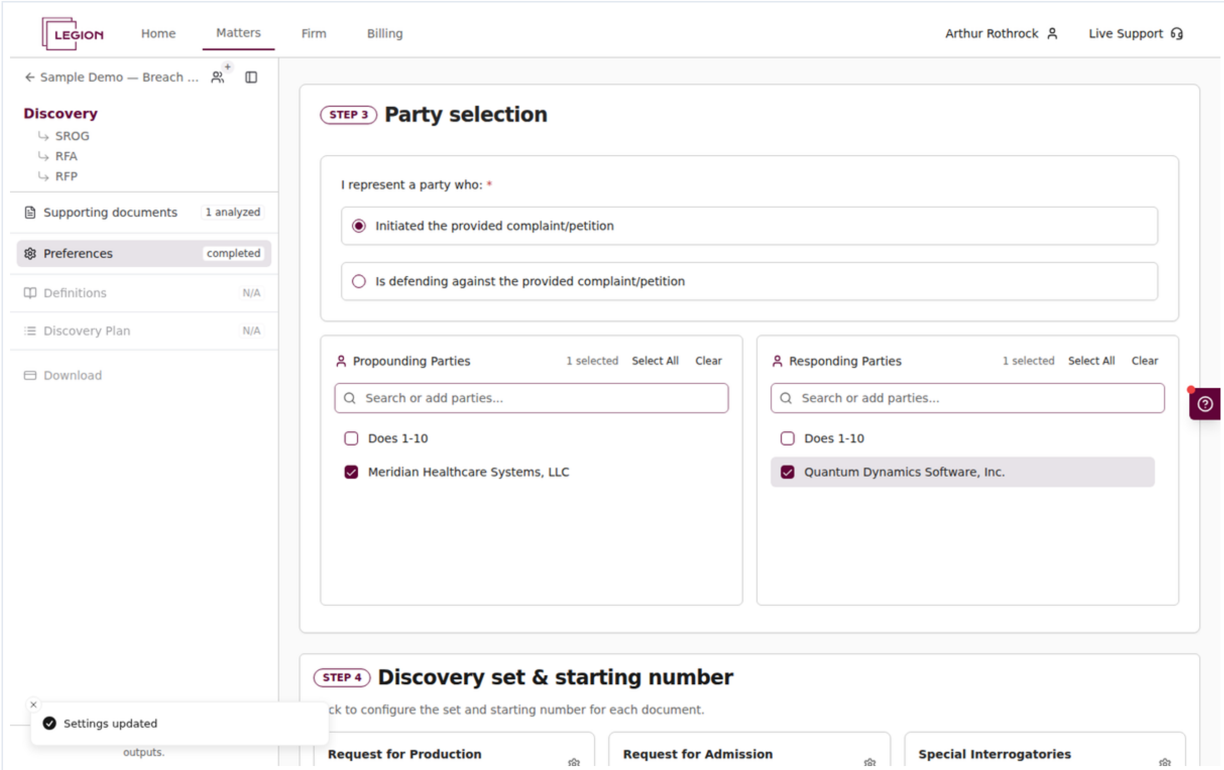

Step 2 — you represent the plaintiff; Meridian asks, Quantum answers.

Step 3 Generate the sets and check coverage with Centurion

Legion drafts every request across all three types; the tabs at the top switch between them. Centurion sits on the right. Click Gap analysis and it maps your requests to each element of every claim and flags what's missing. Apply the suggestions you agree with; nothing changes until you do.

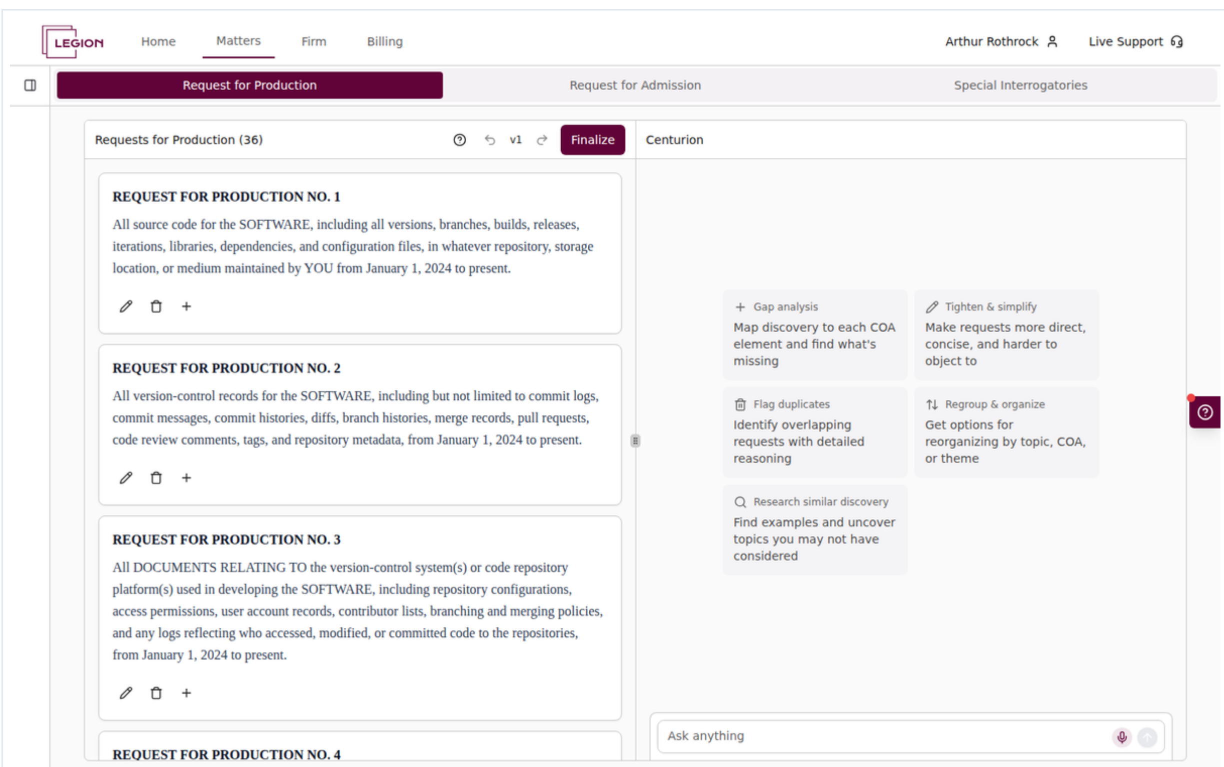

Step 3 — the generated requests, with Centurion's checks on the right.

Step 4 Draft a set, then make one change

When a set looks right, click Finalize to turn it into a formatted document. Read it, then try a revision: type an instruction in the feedback bar and press Enter.

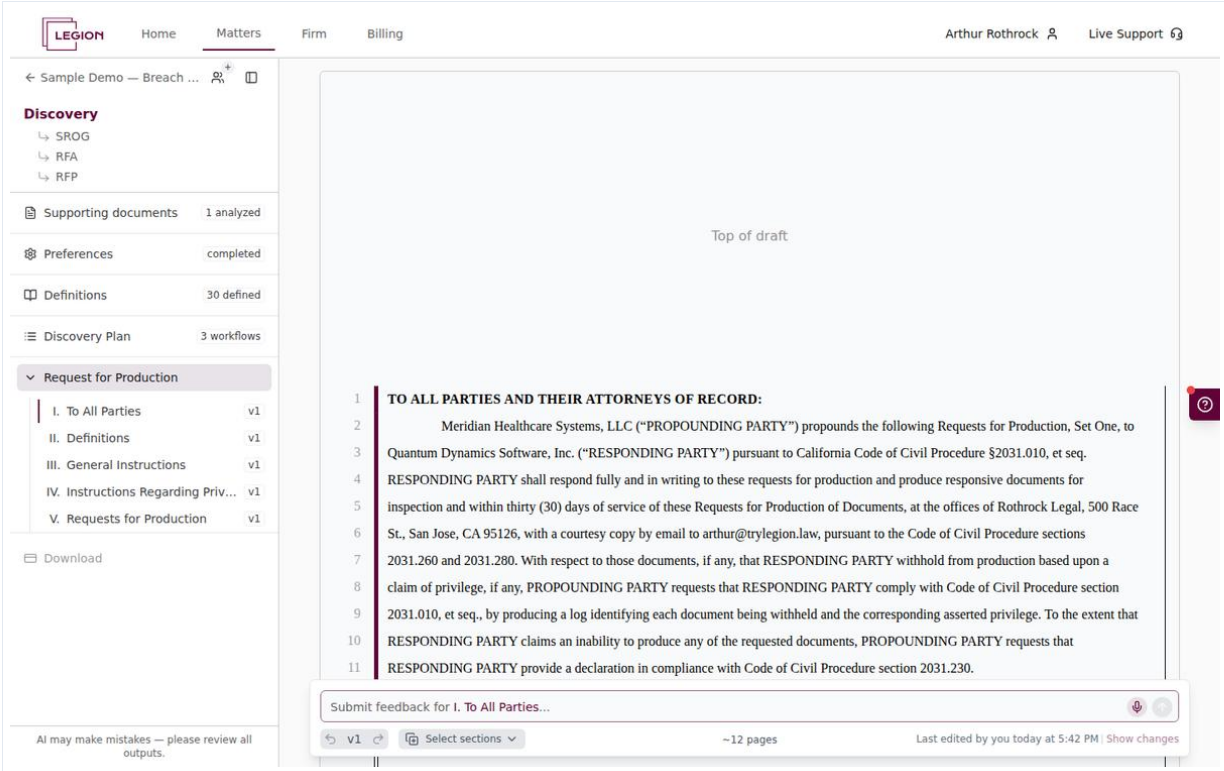

Step 4 — a finalized set, with the feedback bar for revisions.

A REVISION YOU CAN TRY

`Refer to the responding party as "YOU" throughout, and keep each request complete in itself.`

Step 5 Finalize and download

Download each finished set as a Word (.docx) file, ready to finish on your caption and serve.

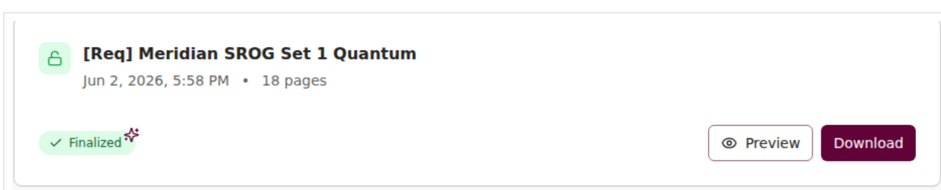

Step 5 — each set ready to download as a Word document.

Done for you: California caps interrogatories and requests for admission at 35 unless counsel signs a declaration of necessity. If a set passes 35, Legion writes that declaration for you and lists it as its own document.

Before you serve

- ☐ Confirm the propounding and responding parties are right before anything else.
- ☐ Centurion suggests; you decide. Don't let "comprehensive" become "everything" — each request is one more objection to fight.

☐ If a counsel declaration was generated, it is a signed court filing — confirm the count, fill the blanks, and add your case number.

Contention Interrogatories

Make the other side commit, in writing, to the facts behind its pleading

A contention interrogatory asks the responding party to state all facts, name all witnesses, and identify all documents that support a specific contention in their pleading. Legion reads the complaint, lets you pick which contentions to probe, and drafts three interrogatories for each one. Follow the five steps below.

You are working in the sample matter Meridian Healthcare Systems, LLC v. Quantum Dynamics Software, Inc. Here you sit on the defense side: you represent Quantum, and you are sending contention interrogatories to the plaintiff, Meridian, demanding the facts behind its complaint. The complaint is already in the matter.

Step 1 Choose the pleading

Open the matter, click Create, and choose Contention Interrogatories (CROG). On the first screen, tick Meridian's complaint as the pleading the interrogatories are built from, then continue.

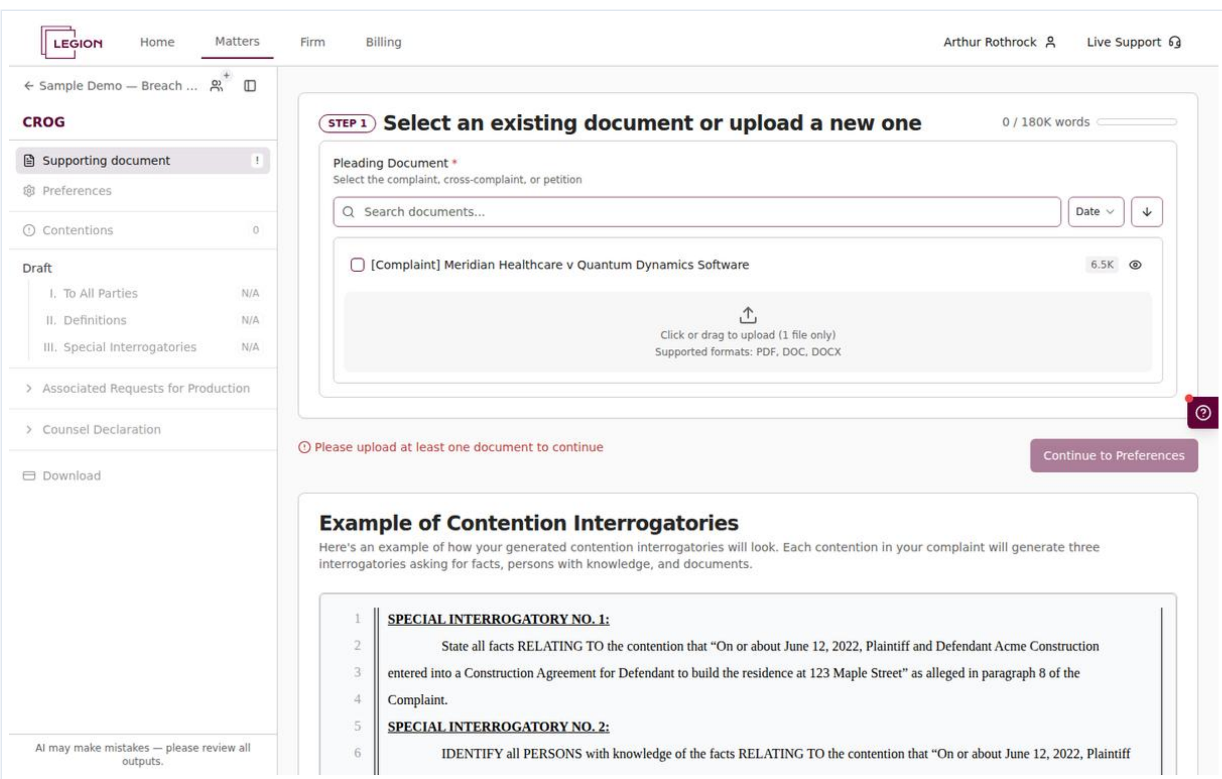

Step 1 — select the complaint already in the matter.

Step 2 Set who is asking whom

This is the one setting that matters most. As Quantum's counsel, set Quantum Dynamics Software as the propounding party (the side you represent) and Meridian Healthcare as the responding party (the side that must answer). Leave the definition style on Separate.

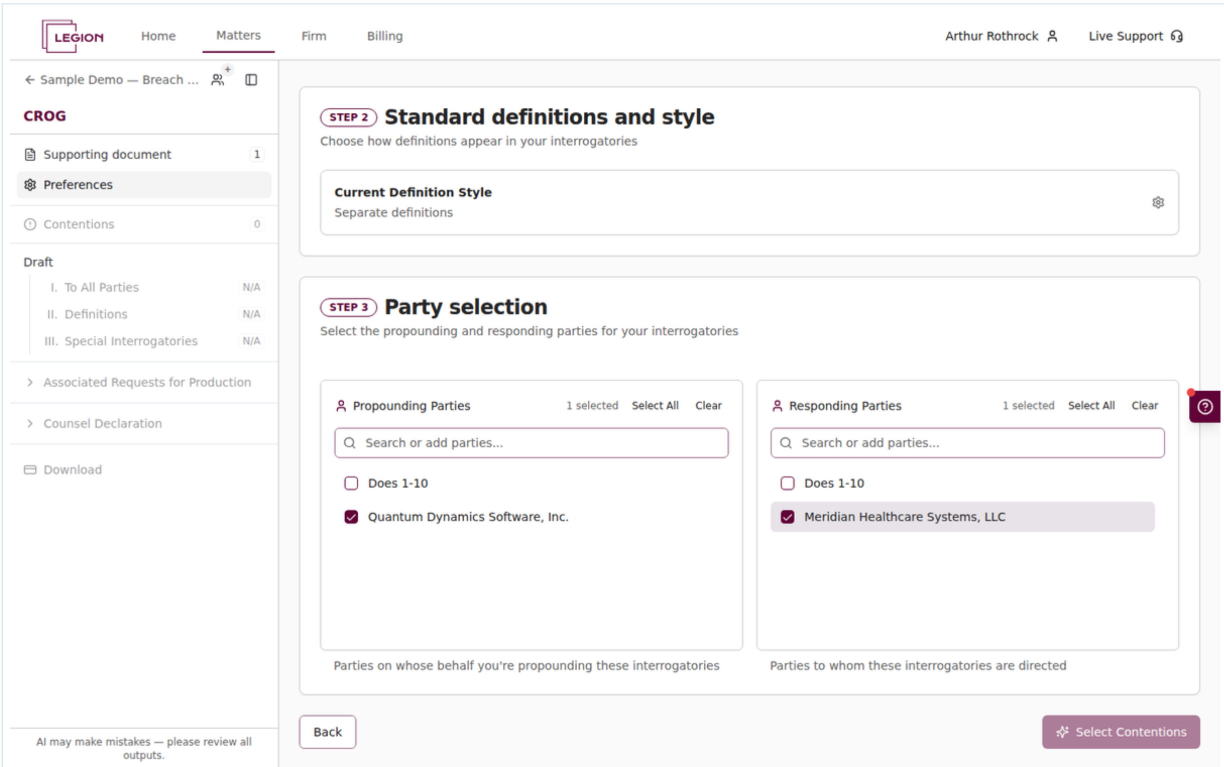

Step 2 — propounding party on the left, responding party on the right.

Step 3 Pick the contentions

Legion highlights the contention sentences in Meridian's complaint. Click the ones worth probing — the load-bearing accusations you want Meridian to commit to. Each becomes a card on the left. When you're done, click Finalize.

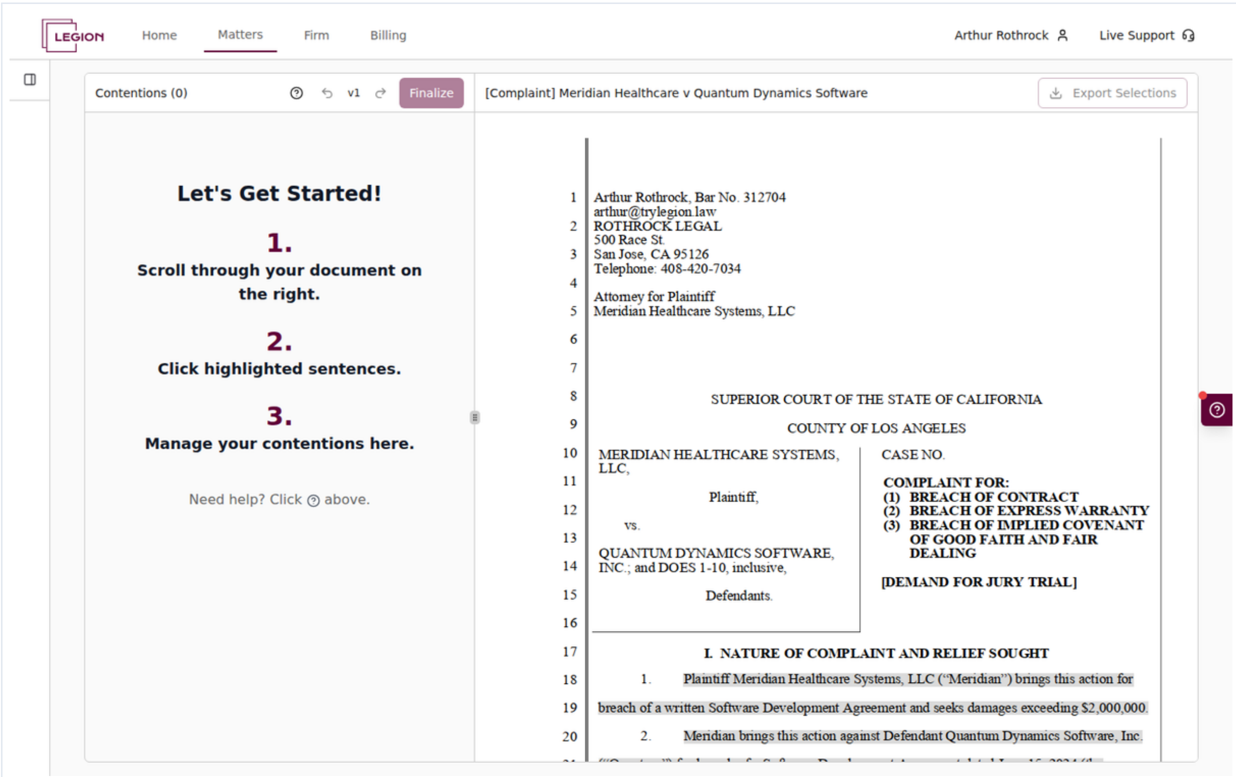

Step 3 — click a highlighted sentence to turn it into a contention.

WHICH CONTENTIONS TO PROBE

`Good ones to pick here: that Quantum missed the Go-Live Date; that it delivered a defective beta; that it failed to implement the required security protocols; and that Meridian seeks damages exceeding $2,000,000.`

Step 4 Draft, then make one change

Legion drafts the full set — three interrogatories per contention — in proper format. Read it, then try a revision: type an instruction in the feedback bar at the bottom and press Enter. Legion rewrites that section and saves a new version.

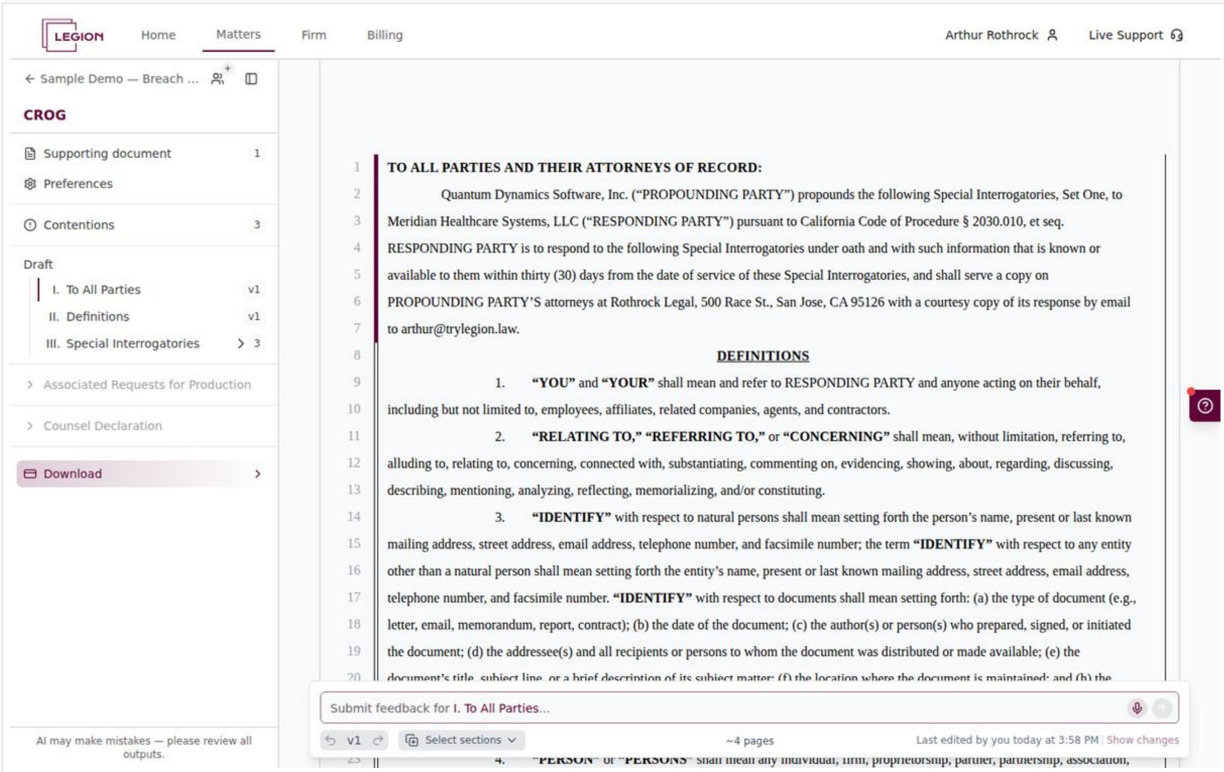

Step 4 — the drafted interrogatories, with the feedback bar at the bottom.

Done for you: Legion also drafts the matching Requests for Production automatically, and adds a short counsel declaration if your set passes 35 interrogatories. You'll find both in the left sidebar.

Step 5 Finalize and download

Click Download to get each finished document as a Word (.docx) file, ready to finish on your caption and serve. If a counsel declaration was generated, fill in its one blank — how many interrogatories you had previously served on this party — before you sign it.

Step 5 — your interrogatories (and the requests that travel with them), ready to download.

Before you serve

☐ Confirm the propounding and responding parties are right — the whole set is addressed based on that choice.

☐ Read the associated Requests for Production; they are a separate served document.

☐ If a counsel declaration was generated, it is a signed court filing — fill the blank, confirm the count, and add your case number.

The MPA Builder

Turn the pleadings and cases in your matter into a drafted Memorandum of Points and Authorities

The MPA Builder does one thing first that sets it apart: it proposes an argument outline you approve before it writes a word of brief. You lock in the structure, Legion drafts each section, and you control every paragraph and citation. Follow the four steps below.

You are working in the sample matter Meridian Healthcare Systems, LLC v. Quantum Dynamics Software, Inc. Here you are defense counsel for Quantum, writing a brief in support of a demurrer to Meridian's tort claims (negligent misrepresentation and professional negligence), which are barred by the economic loss rule. The complaint and the key cases are already in the matter.

> **Important:** Legion drafts only from what you give it. It does not search the web, and it does not check your citations*. Verifying every case and quotation against the original source is on you.
> *Coming soon!

Step 1 Add your documents and instructions

Open the MPA Builder in the matter. Select the complaint and the cases already loaded — Blank v. Kirwan, Barquis, Sheen v. Wells Fargo, and M & L Financial. Then write a short instruction telling Legion what the brief argues, and click Generate Outline.

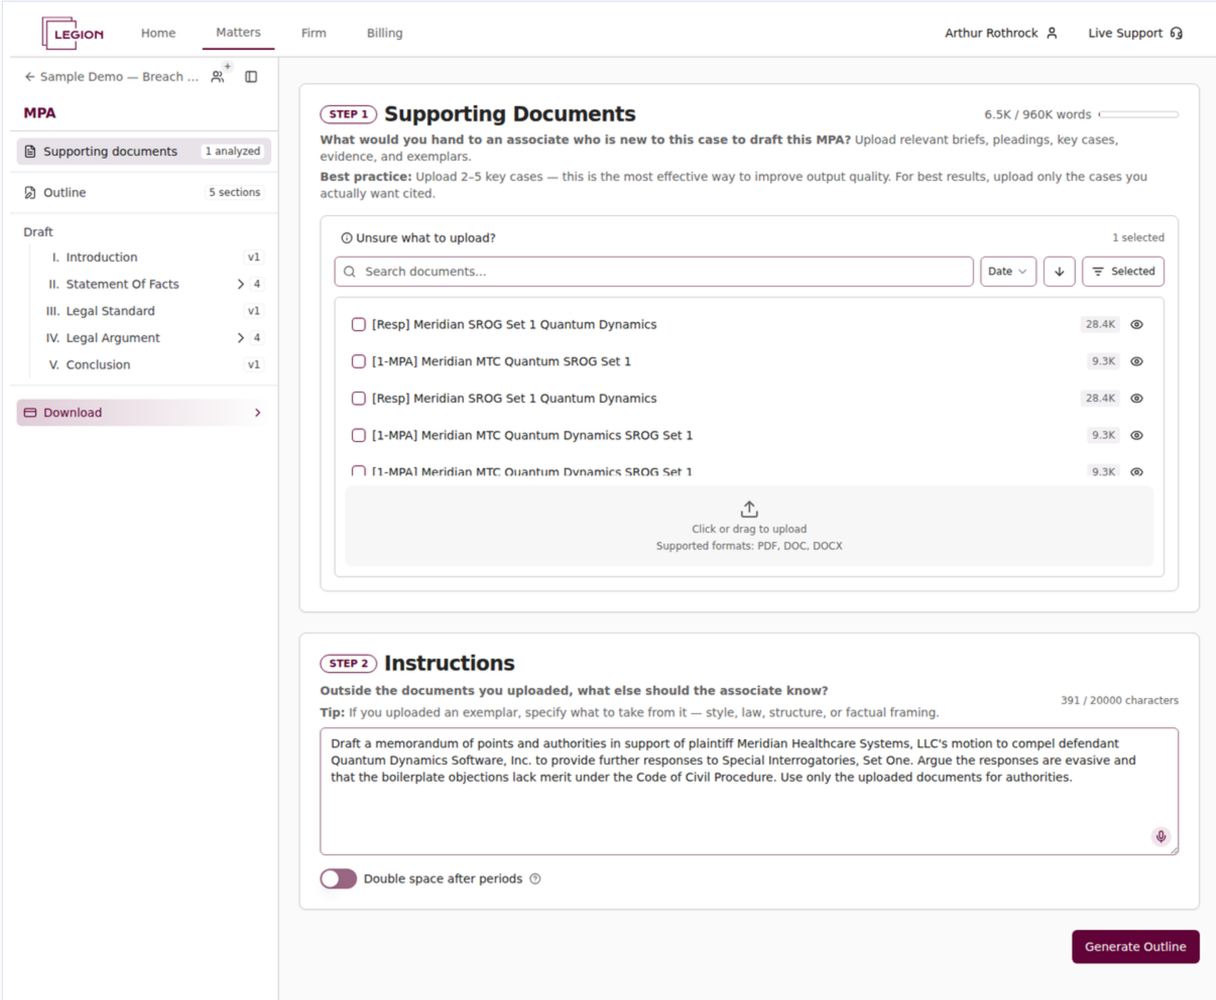

Step 1 — select your documents, then tell Legion what the brief is for.

A SHORT INSTRUCTION YOU CAN TYPE

```
This brief supports Quantum's demurrer to Meridian's tort claims (negligent
misrepresentation and professional negligence). Argue that the economic loss rule
bars them: Meridian's injury is the loss of its contract bargain, not a tort.
Open
with the demurrer standard from Blank v. Kirwan. Cite ONLY the cases I uploaded.
```

Step 2 Approve the outline

Before drafting, Legion proposes the structure of your argument — headings and sub-points. This is the most important step: get it right and the draft follows. Edit any line, or type a request in the feedback box to revise the whole outline. When the structure is right, click Proceed to draft.

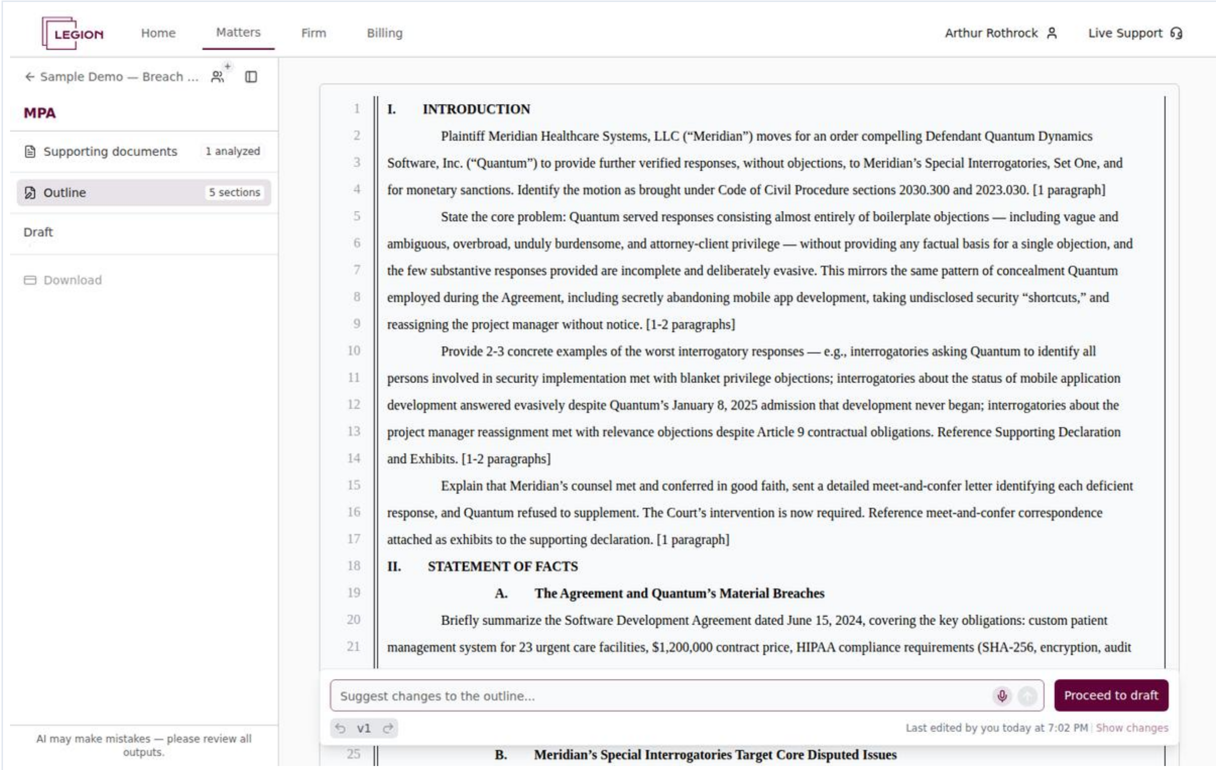

Step 2 — the argument outline. Reorder and refine it before any drafting.

Step 3 Draft the brief, then make one change

Legion drafts each section from your approved outline. Read it, then revise in plain English in the feedback box at the bottom.

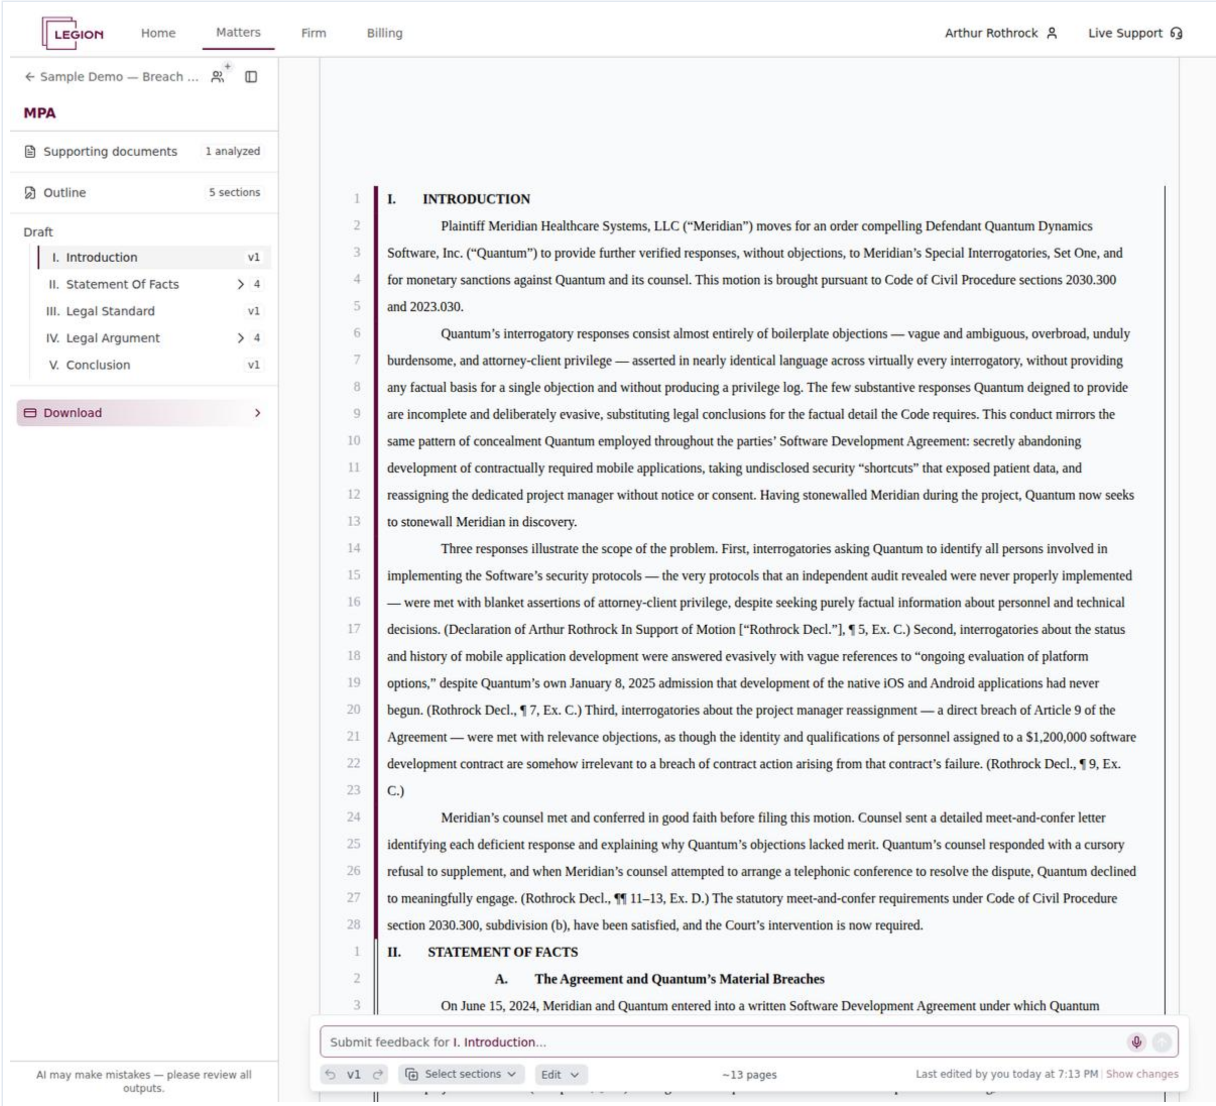

Step 3 — the drafted brief, with the feedback box and the Plan / Edit toggle.

A REVISION YOU CAN TRY

`Tie each alleged misrepresentation to the contractual promise it duplicates, and cut the repetition in the second paragraph.`

Step 4 Finalize and download

When every section is done, click Download in the sidebar to get the brief as a Word (.docx) file, ready to finish on your caption.

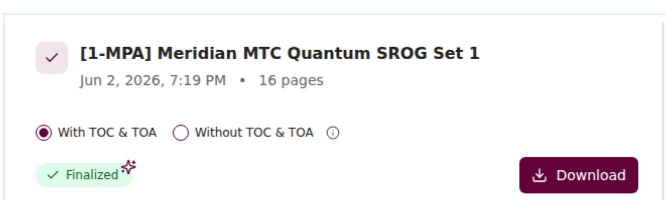

Step 4 — your finished brief, ready to download as a Word document.

Before you file

☐ Read every citation and quotation against the original source. Legion does not check cites, and they can be wrong.

☐ Legion cites only what you uploaded — a fast self-check is to ask it to flag any paragraph relying on a case you did not upload.

☐ You are the signer. The law, the quotes, and the Bluebook form are your responsibility.

The Product Counsel Suite

Four AI Projects Built for Carlos Rosario, Senior Lead Product Counsel at Confluent

---

This worksheet adapts the FRAME method and the projects approach for in-house product counsel work. It was drafted from your public profile and Confluent's public materials, so a few details are bracketed for you to confirm or correct. The shape of the job it assumes: a generalist product counsel at a data streaming platform company, reviewing features and launches, counseling engineers and PMs, handling open source and trademark questions, and backstopping enterprise customer negotiations. Adjust anything that misses.

If This Isn't Quite Your Job, Make It Yours

Maybe your responsibilities at Confluent look different from what's assumed here. Maybe you're another in-house lawyer in the room, doing transactional, commercial, employment, or privacy work instead of product counseling. Either way, don't adapt this worksheet by hand. Have Claude rebuild it around you in a few minutes: open Claude (claude.ai), upload this document, paste the prompt below, and fill in the brackets, especially the responsibilities list. The more specific your responsibilities, the better the projects fit.

```
I've uploaded a worksheet of four AI projects built for a product counsel at a
data streaming company. My role is different, so rebuild it for me.

I'm [NAME], [TITLE] at [COMPANY], which [WHAT THE COMPANY DOES IN ONE SENTENCE].

My main responsibilities, roughly in order of how much of my week they eat: [LIST
3 TO 6, E.G., NEGOTIATING CUSTOMER AND VENDOR CONTRACTS, REVIEWING MARKETING
MATERIALS, SUPPORTING THE SALES TEAM ON DEAL TERMS, EMPLOYMENT QUESTIONS, PRIVACY
REQUESTS].

The recurring tasks I dread most: [ONE OR TWO, E.G., TURNING 40-COMMENT REDLINES,
ANSWERING THE SAME NDA QUESTIONS].

Rebuild the worksheet as follows: keep the overall structure, the FRAME walk-
throughs, the prompt-box format, the one-week sprint, and the ground rules.
Replace the four projects with ones matched to my actual responsibilities, with
the dreaded tasks getting priority. Rewrite the foundation block around me and my
company. Adapt the confidentiality section to my company's situation.

Ask me up to three questions before drafting if anything is unclear. Then return
a new Word document titled "The [MY ROLE] Suite — [MY NAME]".
```

FRAME in one breath: **F**riction is your trigger. **R**each out and ask the AI how. **A**sk well, the CASE way (Context, Ask, Specs, Examples, plus Search when research matters). **M**ake it better through rounds, not restarts. **E**mbed it as a project so you never rebuild it.

The in-house confidentiality gate comes first. Everything below assumes you are working in Confluent's approved AI tooling under its internal AI policy. Product roadmaps, unreleased feature details, customer names, and security findings are company confidential; what you paste into an unapproved tool leaves the building. One more wrinkle unique to in-house life: AI drafts are not attorney work product by default, and mixed business-and-legal communications already sit on shaky privilege ground. Treat project outputs as inputs to your advice, never as the advice itself.

Your Foundation Block

Paste this at the top of every project's instructions after confirming the bracketed parts:

```
About me: I'm Carlos Rosario, Senior Lead Product Counsel at Confluent, the data
streaming platform company built around Apache Kafka and Apache Flink. I operate
as a generalist: product counseling, [PRIVACY / OPEN SOURCE / MARKETING REVIEW /
COMMERCIAL SUPPORT — EDIT TO MATCH].

My background: engineering undergrad, then IP litigation at a large firm, then
in-house product counsel roles at autonomous vehicle and drone delivery companies
before Confluent. I'm comfortable with technical material; do not dumb down
architecture, but do not assume I know this quarter's roadmap either.

My clients are engineers, PMs, and go-to-market teams. They want speed and a
clear path to ship. My job is to find that path while protecting the company, and
to say no clearly on the rare occasions no is the answer.

Confluent context you should assume: products include Confluent Cloud, Confluent
Platform, and stream processing with Flink; customers include enterprises in
regulated industries; the company operates globally; open source is in the
company's DNA and its licensing posture is a live issue.

Standing rules: never invent facts, law, license terms, or contract language.
Flag uncertainty instead of guessing. When my question is ambiguous, ask me
before drafting. California attorney; assume US law unless I say otherwise, and
flag when a question is jurisdiction-sensitive.
```

Project 1 — The Launch Review Partner

A first-pass issue-spotter for PRDs, feature specs, and launch plans. It reads the way product counsel reads: what data does this touch, what are we promising, what could a regulator, customer, or plaintiff make of it.

F Friction is your trigger

A 14-page PRD lands in your queue with a review deadline of yesterday, or a PM pings you with "quick legal check?" on a feature that is anything but quick. The dread of being the bottleneck is the signal.

R Reach out and ask it how

```
I review product launches for legal issues. Here's a sample PRD (sanitized).
Interview me: what does a thorough product counsel review cover at a company like
mine, and what should I give you each time to make your first pass useful?
```

A Ask well: the project instructions

```
You are my first-pass reviewer for product documents: PRDs, feature specs, launch
briefs, and GTM plans.

For each document, produce an issue-spotting memo organized by risk area: (1)
data and privacy: what personal or customer data the feature collects, processes,
stores, or exposes, and in what role (controller vs. processor); (2) commitments:
anything that creates or conflicts with promises in our terms of service, DPAs,
SLAs, security documentation, or certifications; (3) open source and IP: new
dependencies, license obligations, use of third-party models or APIs; (4)
marketing and claims: performance numbers, benchmark comparisons,
"most/best/only" language, and competitor references that need substantiation;
(5) regulated-customer impact: anything that changes the analysis for financial
services, healthcare, or public sector customers; (6) AI-specific issues where
the feature involves AI: training data provenance, customer data use, output
liability, transparency obligations.

Rank issues as ship-blockers, fix-before-GA, or monitor. For every issue, propose
the smallest change that resolves it. My job is finding the path to ship, so a
flag without a fix is half an answer.

End with the five questions I should ask the PM that the document does not
answer.

You are spotting issues, not giving legal advice. Cite no law, regulation, or
contract language unless I have provided it. Where the answer depends on a
document you have not seen (our MSA, our DPA, a specific certification), say
which document and what to check.
```

A model first prompt:

```
CONTEXT: Attached is the PRD for a feature that lets customers route topic data
to a third-party analytics service, plus our standard DPA and the relevant docs
page.

ASK: First-pass review per your instructions.

SPECS: One page. Ship-blockers first. Assume the customer data in scope may
include personal data; we are a processor.

EXAMPLES: Attached is a past review memo of mine so you can see the format and
depth I want.
```

M Make it better

- "Issue 2 is the real one. Draft the three follow-up questions for the eng lead, in plain English, no legalese."
- "You flagged the benchmark claim. Now draft the substantiation checklist marketing needs to complete before we can use it."
- "Re-run the review assuming the feature is GA in the EU on day one. What changes?"

E Embed it

Name it: "Launch Review." **Upload:** your public terms, DPA, and security documentation; one or two of your past review memos as format examples; and your internal launch checklist if one exists. Refresh the PRD per review; the instructions stay put.

Project 2 — The Engineer Whisperer

Turns legal analysis into answers engineers and PMs will actually read: short, technical where it helps, and always ending with what to do.

F Friction is your trigger

You know the answer; writing it for the audience is the chore. The Slack thread is waiting, the PM wants a yes or no, and the honest answer is three paragraphs of "it depends" that nobody will read. That gap between your analysis and their attention span is the friction.

R Reach out and ask it how

```
I counsel engineers and PMs. My answers are legally right but too long and too
hedged for the audience. Here are two recent answers I wrote. How would you help
me convert analysis like this into guidance the team will read and follow?
```

A Ask well: the project instructions

```
You convert my legal analysis into guidance for engineers and PMs.

Format: lead with the answer (yes, no, or yes-if), then the reasoning in three
sentences or fewer, then exactly what the team needs to do, as concrete steps
tied to their workflow (the PRD, the design doc, the config, the docs page).

Tone: a technically fluent colleague, not a law firm memo. Engineers respect
precision; never trade accuracy for simplicity. If the simple version would be
wrong, say "it's more complicated" and give the shortest correct version.

Preserve every substantive caveat I include, but convert it from hedge to
instruction: not "this may raise issues under applicable law" but "if the data
includes EU personal data, stop and come back to me."

Produce two versions: a Slack-length reply and a paragraph for the design doc or
wiki.

Never alter my legal conclusions. If my draft is ambiguous on the conclusion, ask
me rather than picking one.
```

A model first prompt:

```
CONTEXT: An eng team wants to log sample payloads from customer topics to debug a
connector. Below is my analysis: it's prohibited by default under our DPA, with a
narrow path involving customer consent and redaction.

ASK: Convert per your instructions.

SPECS: The Slack version must survive being read in eight seconds. The wiki
version should be reusable for the next team that asks.

EXAMPLES: Here are two past answers of mine that landed well; match their
bluntness.
```

M Make it better

- "The eng lead read it and asked 'so can we ship Friday?' The answer didn't land. Make the conditional clearer."
- "Turn this answer into a decision tree the team can self-serve next time, with the 'call legal' branches marked."
- "Draft the one-slide version for the product review meeting."

E Embed it

Name it: "Engineer Whisperer." **Upload:** your best past counseling answers as voice examples (sanitized), plus any internal FAQ or wiki pages you maintain. Borrow the voice-quote trick from the correspondence project: have it pull a dozen short quotes that capture how you explain things, and keep them in the project.

Project 3 — The Open Source Sentinel

A triage assistant for the licensing questions that come with building on Apache Kafka and Flink: dependency reviews, license compatibility, contribution questions, and use of third-party marks.

F Friction is your trigger

An engineer asks "can we use this library?" and the honest workflow is twenty minutes of reading a license you've read forty times before, plus checking how the dependency is actually used. Multiply by every sprint. This project does the first fifteen minutes.

R Reach out and ask it how

```
I handle open source questions at a company whose products are built around
Apache-licensed projects and whose own licensing posture includes both open
source and source-available components. Walk me through how you'd triage a new-
dependency request, and tell me what context files would make your triage
reliable.
```

A Ask well: the project instructions

```
You triage open source licensing questions. You are a first-pass analyst, not the
decision-maker.

For a new dependency request, work through: what license the component claims
(and tell me to verify the actual LICENSE file, not the README); how we would use
it (static linking, dynamic linking, modification, distribution, SaaS-only);
whether that use triggers obligations (notice, attribution, source disclosure,
patent provisions); and how it interacts with the licenses of our own products
and the upstream Apache projects.

Classify each request: green (standard permissive license, standard use, document
and go), yellow (obligations exist, here's the compliance step), red (copyleft
conflict, license ambiguity, or anything touching our proprietary components —
goes to me with a summary of exactly what's unresolved).

For trademark questions about third-party project names and marks, identify the
applicable trademark policy to check rather than guessing at what it permits.

Hard rules: quote license language only from license text I provide or that you
retrieve and identify by exact version. Never paraphrase a license obligation
without flagging that paraphrase. If the component's license is unusual or
custom, that is automatically red.

Maintain a running log format I can paste into our tracking system: component,
version, license, use, classification, obligations, decision.
```

A model first prompt:

```
CONTEXT: Eng wants to add [LIBRARY] to a Confluent Platform component. Attached:
the library's LICENSE file, its NOTICE file, and a paragraph from the eng lead on
how it will be used.
```

```
ASK: Triage per your instructions.

SPECS: Classification first. If yellow, the compliance step must be specific
enough to paste into the ticket.

SEARCH: Confirm the current license version for this library's latest release
before you classify.
```

M Make it better

- "This one's yellow. Draft the NOTICE file entry and the attribution language for the docs."
- "The eng lead says they'll only use it server-side. Re-triage with that assumption and tell me whether it changes the classification."
- "Summarize this month's log: classifications by color, and the recurring patterns I should turn into self-serve guidance."

E Embed it

Name it: "OSS Sentinel." **Upload:** your internal open source policy, your approved/denied component list, the license texts you see most, and your products' own license files. The classification rules in the instructions should mirror your actual policy; edit them until they do.

Project 4 — The Customer's Counsel

The sparring project, in-house edition. It plays the enterprise customer's lawyer in a negotiation: redlining your DPA, pushing on AI terms, demanding audit rights, so you walk into the call having already heard the arguments.

F Friction is your trigger

A negotiation call is on the calendar with a customer's counsel known to be aggressive, or a redline arrives with forty comments and a deal team asking how fast you can turn it. Preparing alone, you only think of the arguments you already know.

R Reach out and ask it how

```
You're going to play counsel for a large [REGULATED INDUSTRY] enterprise
negotiating our cloud agreement and DPA. What do you need to play them well:
their likely regulatory drivers, their published vendor requirements, our
standard terms?
```

A Ask well: the project instructions

```
You play counsel for an enterprise customer negotiating against my company's
standard agreements. Stay in character as a sophisticated, demanding, but
commercially rational lawyer.

MODE 1, THE REDLINE. Given our standard terms and the customer's profile
(industry, jurisdiction, regulatory drivers), produce the redline their counsel
would send: the clauses they'd strike, the language they'd insert, and the one-
line justification they'd give for each. Mark which asks are real requirements
driven by their regulators and which are wish-list items they'll trade away.

MODE 2, THE CALL. When I say "get on the call," argue your top three positions
against me live, one at a time. Push back on my responses the way real counsel
would. When I say "end call," step out of character and score it: which of my
arguments worked, where I conceded too fast, what I should put in the follow-up
email, and what their walk-away points likely are.

Ground rules: argue only from the documents and facts I provide plus reasonable
industry-standard positions, and label anything that is an assumption about this
specific customer. You know nothing about the real customer's actual
instructions; you are a rehearsal, not intelligence.
```

A model first prompt:

```
CONTEXT: Attached are our standard cloud terms and DPA, plus the customer's
profile: EU-headquartered bank, DORA-driven vendor requirements, prior deal fell
apart over audit rights.

ASK: Mode 1. Send me their redline.

SPECS: Top ten asks only, ranked by how hard they'll fight. Mark regulatory must-
haves versus tradeables.
```

```
EXAMPLES: Attached is a (sanitized) redline from a similar customer last year;
calibrate aggressiveness to that.
```

M Make it better

- "Get on the call. Audit rights first."
- "You folded on the liability cap too easily. Real bank counsel wouldn't. Run it again, harder."
- "End call. Now switch sides and draft my follow-up email memorializing where we landed."

E Embed it

Name it: "Customer's Counsel." **Upload:** your standard terms, DPA, and security exhibit (public versions where possible), plus sanitized redlines from past negotiations sorted by customer type. The past redlines are what make the simulation realistic instead of generic.

A One-Week Sprint to Stand These Up

1. **Monday:** confirm the foundation block, fix the brackets, and clear the confidentiality gate with whoever owns Confluent's internal AI policy.
2. **Tuesday:** build Launch Review and run it against a PRD you already reviewed by hand. Compare its memo to yours; tune the instructions on the gaps.
3. **Wednesday:** build Engineer Whisperer using two past answers as voice samples. Use it on the next real Slack question.
4. **Thursday:** build OSS Sentinel and re-run last month's dependency requests through it. Check its classifications against your actual calls.
5. **Friday:** build Customer's Counsel and rehearse your next negotiation. Then ask each project: "Rewrite your own instructions to fix what you got wrong this week."

Ground Rules for In-House Use

☐ Approved tooling only, under Confluent's internal AI policy. Roadmap, customer, and security information never goes into unapproved tools.

☐ Project outputs are inputs to your advice. The legal call, and anything that reads like one, comes from you.

☐ License terms, contract language, and regulatory requirements get verified against the source document every time, regardless of how confident the output sounds.

☐ Privilege check: AI drafts and chats are not privileged by default. Keep the analysis-in-progress in your head or your notes, and label your actual advice appropriately.

☐ The simulated customer counsel knows nothing about the real customer. Rehearsal, not intelligence.

☐ When a project's output surprises you on a question that matters, that's a conversation with a colleague or outside counsel, not another prompt.

Why these four: product counsel work is high-volume triage punctuated by judgment calls. The projects take the triage: the first-pass review, the translation, the license read, the rehearsal, so the hours you have go to the judgment. That's the part they can't do, and the part the company actually pays for.

The Family Law Suite

Four AI Projects for Family Law Attorneys, Built with the FRAME Method

---

Family law runs on volume and emotion: declarations due on RFO timelines, clients who send fourteen-paragraph emails at midnight, disclosures that arrive as a shoebox of statements, and hearings in front of bench officers juggling forty matters before lunch. The four projects in this packet take the volume so your hours go to judgment and strategy. Each is walked through the FRAME method so you can see how to build the next one yourself.

FRAME in one breath: **F**riction is your trigger. **R**each out and ask the AI how. **A**sk well, the CASE way (Context, Ask, Specs, Examples, plus Search when research matters). **M**ake it better through rounds, not restarts. **E**mbed it as a project so you never rebuild it.

The Confidentiality Gate, Family Law Edition

Family law files are the most sensitive in civil practice: children's information, abuse allegations, mental health records, and the complete financial life of two households. Before any project below touches real matter material:

- ☐ Firm-approved AI account only, on a plan where inputs are not used to train models. Never a personal free account.
- ☐ No real names. Refer to parties as Petitioner/Respondent or by initials, children by initials and ages, and strip case numbers before pasting anything in.
- ☐ Extra caution with domestic violence and child abuse content. Summarizing it for a declaration is one thing; uploading a forensic evaluation or a CPS report needs a hard look at your firm's policy first, and the answer may be no.
- ☐ Financial documents get sanitized: account numbers redacted, names replaced, before upload.
- ☐ Check your state bar's guidance on generative AI. California attorneys: the State Bar's Practical Guidance on generative AI addresses confidentiality, supervision, and candor directly.

Your Foundation Block

Paste this at the top of every project's instructions, filled in:

```
About me: I'm [NAME], a family law [PARTNER/ASSOCIATE/SOLO] at [FIRM] in
[COUNTY], California. My caseload is [DISSOLUTIONS / CUSTODY DISPUTES / SUPPORT
MODIFICATIONS / DVRO MATTERS / HIGH-ASSET CASES — EDIT].

My courts: [COUNTY] Superior Court, family division. My matters move on RFO and
trial timelines; deadlines are short and the judges read fast.

My clients are going through the worst year of their lives. They need clear,
calm, honest communication, and they pay hourly, so efficiency is kindness.

Standing rules: never invent facts, dates, court rules, statutes, or case law.
Work only from materials I provide. Refer to parties as Petitioner and Respondent
and to children by initials unless I say otherwise. Flag uncertainty instead of
guessing, and ask me before drafting if my request is ambiguous.

Never compute guideline child or spousal support; that comes from certified
calculator software. You may organize the inputs and explain what's missing.
```

My caseload mix: ____

My county and department quirks worth telling the AI:

____

The two tasks I dread most each week:

____

Project 1 — The Declaration Builder

Turns interview notes, client timelines, and text-message exhibits into a clean, fact-driven declaration a family court judge can absorb in four minutes.

F Friction is your trigger

The RFO is due, your notes are three pages of client narrative looping through five years of grievances, and somewhere in there are the eight facts that matter. Untangling narrative into declaration is the dread; hand it over.

R Reach out and ask it how

```
I draft declarations supporting requests for orders in family court. Here are my
sanitized notes from a client interview. Before drafting anything, tell me: what
do you need from me to turn notes like these into a declaration, and what should
I have the client gather?
```

A Ask well: the project instructions

```
You help me draft declarations for family court filings: RFOs, responsive
declarations, and supporting declarations.

Structure: first person, the declarant's own voice, numbered paragraphs,
chronological within topics. Open with who the declarant is and the relief
sought, then the facts, then a short closing paragraph.

Facts only. No argument, no characterizations of the other party's motives, no
adjectives doing the work evidence should do. "He arrived 45 minutes late to six
of the last eight exchanges" beats "He constantly disrespects the schedule." If
my notes contain a characterization, convert it to the underlying facts or flag
that I need facts to support it.

Every fact tied to its source: the declarant's observation, a specific text or
email (identify it as an exhibit), a record, or a date. Where my notes give no
source, mark the sentence [SOURCE?].

Length discipline: family court judges read fast. Default to three pages or
fewer. If the material genuinely needs more, tell me what to cut to get there
anyway, then give me the longer version too.

Match the declarant's voice: plain words, the way this client actually talks, not
lawyer language. Never put words in the declarant's mouth; everything traces to
my notes.

Never invent facts, dates, or events. The notes are the universe.
```

A model first prompt:

```
CONTEXT: RFO to modify the parenting schedule. Attached: my sanitized interview
notes, the current order's schedule terms, and a log of exchange times from the
co-parenting app.
```

```
ASK: Draft the client's supporting declaration.

SPECS: Three pages max. Lead with the schedule disruptions, then the effect on
the child (initials only). Mark every unsourced sentence [SOURCE?].

EXAMPLES: Attached is a declaration of mine the court found persuasive; match its
tone and paragraph rhythm.
```

M Make it better

- "Paragraphs 6 through 9 read like argument. Convert each to bare facts and tell me which exhibits I need."
- "The client is a soft-spoken nurse and this sounds like a prosecutor. Rewrite in her register."
- "Now draft the responsive declaration the other side will file, so I can see what mine has to survive."

E Embed it

Name it: "Declaration Builder." **Upload:** two or three of your best past declarations (sanitized) as voice and format examples, plus your interview-notes template if you use one. Per matter, add the sanitized notes and the operative orders.

Project 2 — The Client Whisperer

Two jobs in one project: distills the midnight email rant into facts and action items, and drafts your replies, calm, clear, and billable-hour honest.

F Friction is your trigger

A 1,400-word client email arrives, equal parts fear, fury, and four actual questions buried in paragraph nine. You dread both the excavation and the reply. That's two frictions, and this project takes both.

R Reach out and ask it how

```
A large share of my week is reading long, emotional client emails and answering them without pouring fuel on anything. Here are two sanitized examples. How would you help me process these faster while keeping my clients feeling heard?
```

A Ask well: the project instructions

```
You help me manage family law client communications. Two modes.

MODE 1, THE DISTILLER. When I paste a client email or call notes, extract: (1) new facts, with dates; (2) the client's actual questions; (3) action items for me, for staff, and for the client; (4) anything with a deadline or court-rule implication; (5) emotional temperature and any safety concern that needs immediate attention, flagged at the very top. Safety concerns include any mention of violence, threats, self-harm, or risk to children; never bury these.

MODE 2, THE REPLY. Draft my response: acknowledge the emotion in one sentence without amplifying it, answer the questions in plain English, state what happens next and when, and gently redirect anything that isn't a legal question (therapy, co-parenting apps, the other party's new partner) to the right resource. Honest about costs: if the client's request means billable work, the draft says so and asks before proceeding.

Tone: warm, steady, never patronizing. Short paragraphs. No legal jargon without a plain-English translation in the same sentence.

Hard rules: never give the client false comfort about outcomes. Never speculate about the other party's motives. Never include settlement positions or strategy in writing unless I explicitly include them.
```

A model first prompt:

```
CONTEXT: Below is a sanitized client email about the other parent introducing a new partner to the kids, plus three questions about the upcoming mediation.

ASK: Mode 1 first, then Mode 2.

SPECS: Reply under 250 words. The mediation questions get real answers; the new-partner topic gets acknowledged and redirected unless a court order addresses it (none does).
```

```
EXAMPLES: Attached are two past replies clients responded well to; that's my
voice.
```

M Make it better

- "Too clinical. She's scared, not stupid. Warmer open, same substance."
- "Turn the action items into a task list for my paralegal with deadlines attached."
- "This client emails daily. Draft a kind paragraph establishing a weekly update rhythm without making her feel cut off."

E Embed it

Name it: "Client Whisperer." **Upload:** two or three past client emails you're proud of (sanitized), and use the voice-quote trick: "Pull about 12 short quotes that capture how I talk to clients, across calm and crisis," then keep those quotes in the project.

Project 3 — The Settlement Architect

Organizes the financial side of a dissolution: builds the asset and debt picture from disclosures, drafts division scenarios, and hunts for what's missing before you sign off on an MSA.

F Friction is your trigger

Final declarations of disclosure are exchanged, the FL-142 runs to forty line items, and settlement is close. Building the division spreadsheet and stress-testing the deal is hours of careful tedium with malpractice stakes. Tedium with stakes is exactly what to frame.

R Reach out and ask it how

```
I settle dissolution cases. The inputs are each side's Schedule of Assets and
Debts and Income and Expense Declaration, plus statements. Walk me through how
you'd help me organize a community estate for division and check a proposed deal
for holes, and tell me exactly how the documents should be sanitized before I
share them.
```

A Ask well: the project instructions

```
You help me analyze and settle the financial side of dissolution matters. Work
only from the sanitized disclosures, statements, and term sheets I provide.

TASK 1, THE ESTATE PICTURE. Build a table of assets and debts: item, claimed
character (community, separate, disputed), each side's stated value, valuation
date, and source document. Flag every mismatch between the two parties'
disclosures, every asset referenced in one document but missing from another (a
statement showing transfers to an account not disclosed, employment income
suggesting retirement benefits nobody listed), and every item with a stale or
missing value.

TASK 2, DIVISION SCENARIOS. Given a proposed split, calculate each side's net and
the equalization payment, and present two or three alternative allocations
reaching the same net result with different assets. Show arithmetic so I can
check it.

TASK 3, THE HOLE HUNT. Review a draft MSA or term sheet against the estate
picture and list what is unaddressed: items in the disclosures missing from the
agreement, and standard issue categories for me to confirm (tax consequences of
the division, reimbursement and credit claims, treatment of retirement accounts
and whether orders dividing them are needed, debts in joint names, support
interplay, life insurance securing support). Present these as questions for my
judgment, not conclusions.

Hard rules: never compute guideline support; identify inputs and gaps only. Never
opine on what a court would do. Arithmetic is yours; legal characterization is
mine — present characterization questions, don't resolve them. If a number
doesn't trace to a document I gave you, don't use it.
```

A model first prompt:

```
CONTEXT: Attached are both parties' sanitized FL-142s and FL-150s, the last three
statements for the brokerage account, and opposing counsel's proposed term sheet.

ASK: Task 1, then Task 3 against the term sheet.

SPECS: Table format I can paste into Excel. Mismatches and missing items in their
own section at top. Questions ranked by dollar exposure.

EXAMPLES: Attached is my division spreadsheet from a prior matter; mirror those
columns.
```

M Make it better

- "The pension shows up in his FL-142 but not the term sheet. Draft the three questions I send opposing counsel about it."
- "Run a scenario where she keeps the house. What does the equalization look like and what does it depend on?"
- "Convert the open items into my settlement-conference checklist, ordered by what to resolve first."

E Embed it

Name it: "Settlement Architect." **Upload:** your division spreadsheet template, a sanitized model MSA you trust, and your standard settlement checklist. Per matter, the sanitized disclosures. The hole-hunt is only as good as the checklist you teach it, so invest there.

Project 4 — The Bench Officer

The sparring project. A family court judge with a crowded calendar reads your RFO papers, tells you the tentative inclination, and asks the questions you'll hear from the bench, including the uncomfortable ones about your client.

F Friction is your trigger

The hearing is Thursday, you get maybe ten minutes of argument, and the judge will have read the papers the night before at best. You're rehearsing in the car. Rehearse against resistance instead.

R Reach out and ask it how

```
You'll play the family court bench officer hearing my RFO. What helps you do that
realistically: both sides' papers, the operative orders, the local rules, this
department's tentative practice?
```

A Ask well: the project instructions

```
You are an experienced family court bench officer with a heavy calendar. You care
about the children's interests, workable orders, and credibility; you have no
patience for lawyers arguing about the parents' grievances instead of the
requested relief.

When I give you the moving and responding papers, do four things:

First, in two paragraphs, your tentative inclination and why, including what
evidence or concession would move you.

Second, the 10 to 12 questions you would ask from the bench, most likely first.
Include the hard ones for my side: the gaps in my declaration, my client's
conduct, why my proposed order is workable day-to-day, and what I'd accept short
of everything. Include two questions you'd ask opposing counsel that I should be
ready to use.

Third, for each question, one line on the answer that satisfies you and one line
on the answer that hurts me.

Fourth, the practical defects: relief requested that the notice doesn't support,
orders too vague to enforce, exhibits I reference but didn't attach.

Stay in character. Work only from the papers I provide; invent no facts, no
rules, no case law. You decide nothing about the real hearing; you exist to make
me ready for it.
```

A model first prompt:

```
CONTEXT: Attached: my RFO to modify custody, the responsive declaration, my
reply, and the current orders. Department holds 15-minute hearings; no tentatives
published.

ASK: Tentative, bench questions, and defects, per your instructions.
```

```
SPECS: Blunt. The questions about my client's relocation timing should be as hard
as the real ones will be.

EXAMPLES: None needed; the papers are the record.
```

M Make it better

- "Ask me your questions one at a time and grade my answers. Follow up the way you would on the bench."
- "Rule against me and dictate the order from the bench, so I hear the findings I have to head off."
- "My ten minutes: which three points do I lead with, and which do I let the papers carry?"

E Embed it

Name it: "Bench Officer." **Upload:** the current matter's papers and operative orders, your county's local family rules, and any standing orders for the department. Swap the papers each hearing; the judge stays.

Make It Yours

If your practice tilts differently (heavy DVRO work, mostly mediation, high-asset dissolutions, juvenile dependency crossover), don't adapt by hand. Open Claude, upload this packet, and paste:

```
I've uploaded a packet of four AI projects for family law attorneys. Rebuild it for my practice.

I'm [NAME], a family law attorney at [FIRM] in [COUNTY], California. My caseload: [DESCRIBE].

My main responsibilities, by share of my week: [LIST 3 TO 6].

The recurring tasks I dread most: [ONE OR TWO].

Keep the FRAME walk-throughs, the prompt boxes, the confidentiality gate, and the ground rules. Replace the projects with ones matched to my responsibilities, dreaded tasks first. Ask me up to three questions before drafting, then return a new Word document titled "The Family Law Suite — [MY NAME]".
```

Run One Live (Workshop Exercise)

1. **Pick a closed matter** so nothing live is at risk, and sanitize one document from it: interview notes, a client email, or an RFO.
2. **Build one project** from this packet, foundation block first, and run the sanitized document through it.
3. **Compare** the output to what you actually filed or sent. Tell the AI exactly where it fell short and let it revise; that's the M step doing its work.
4. **Embed:** ask "Rewrite your own project instructions so you'd have gotten this right the first time," update the project, and save it.

In a group, pair up for the Bench Officer: trade papers, read each other the AI's bench questions cold, and count who survives more of them. Family lawyers argue for a living; use that.

Ground Rules for Family Law Use

☐ The confidentiality gate on page one is not optional, and it's stricter here than anywhere else in this series. Children's information and abuse allegations get the most protective treatment your firm allows.

☐ Guideline support numbers come from certified calculator software, never from the AI. It organizes inputs; the calculator computes.

☐ Declarations are the client's testimony under penalty of perjury. The client reads every word and confirms every fact before signing; the AI's draft is scaffolding, nothing more.

☐ Court forms are the Judicial Council's. The AI drafts content; the form fields, the formatting, and the filing requirements get checked against the real form and your local rules.

☐ Verify every statute, rule, and case against the source before it appears in anything filed. The anti-invention instructions lower the risk; they don't retire it.

☐ The simulated bench officer predicts questions, not outcomes, and it has never met your judge. Calibrate accordingly.

☐ You are the final reviewer of everything that reaches a client, opposing counsel, or the court.

The payoff: family law rewards the attorney who can absorb a chaotic record fast and stay steady with a frightened client. These projects buy back the absorbing hours: the untangling, the distilling, the spreadsheet. What they hand back is time with the parts of the file, and the person, that need a lawyer.

Drafting Firm Articles with AI

A Hands-On Worksheet for Attorneys, Built on the Berliner Cohen Article Style

---

This worksheet walks you through using an AI assistant (such as Claude or another large language model) to produce a first draft of a client-facing article consistent with the style of your firm's news and insights page. You supply the legal substance; the AI supplies a structured first draft; you remain the author and final reviewer. The examples and templates that follow were built from Berliner Cohen LLP's published articles. The very first section shows you how to regenerate this entire worksheet around your own firm's website.

Time required: 30–45 minutes. **What you need:** a topic in your practice area, access to an AI chat tool approved by the firm, and the authorities you intend to rely on.

Before You Begin — Customize This Worksheet for Your Own Firm

This worksheet was built from sample articles on Berliner Cohen's website, but the same exercise works for any firm. You can have Claude rebuild the entire worksheet around your firm's publications in a few minutes:

1. **Open Claude** (claude.ai or the Claude app) and make sure web search is enabled (it's in the tools/settings menu).
2. **Upload this document** to the chat using the paperclip/attach button.
3. **Paste the prompt below**, replacing the bracketed fields with your firm's information, and send it.
4. **Download your new worksheet.** Claude will return a revised Word document built around your firm's actual articles, formats, voice, and contact conventions, ready to use for Parts 1 through 4.

Customization prompt (copy, fill in, paste into Claude)

```
I've uploaded a workshop worksheet that teaches attorneys to draft client-facing
articles with AI. It was built around the article style of one firm. I want you
to rebuild it for my firm.

My firm: [FIRM NAME]. Our website: [URL]. Our articles/insights page, if I know
it: [URL OR "please find it"].

Please do the following:

1. Visit our site and read 3–5 representative articles or client alerts from our
news/insights page. If we publish in distinct formats (alerts, longer insights,
case-law explainers, newsletters), sample at least one of each.

2. Decode our house style the way Part 1 of the worksheet does: article types,
typical length, audience, voice, jurisdiction focus, heading conventions,
citation style, and how our articles close (contact language, boilerplate,
attorney bylines).

3. Rewrite the worksheet's Style Block and prompt templates (Part 2) so they
reproduce OUR style, including our practice areas, office locations, and closing
conventions, rather than the original firm's.

4. Keep the overall structure: style decode, reusable prompt components, the
hands-on exercise with fill-in fields, and the full attorney review checklist in
Part 4 (keep every checklist item; add any items our jurisdiction or practice mix
requires).

5. If our site has too few substantive articles to decode a style, say so, and
instead build the templates from our practice area pages and ask me a few
questions about the voice we want.

6. Return the result as a new downloadable Word document titled "AI Article
Drafting Worksheet – [FIRM NAME]".
```

Tip: the more specific you are in the bracketed fields, the better the result. If your firm's best writing lives in a newsletter PDF or on a specific attorney's blog, upload or link those too and tell Claude to treat them as the style samples.

Part 1 — The House Style, Decoded

Three recent articles were pulled from berliner.com as samples. Each represents a recurring format:

- **Client Alert**: e.g., "California Employer Alert: February 1st Deadline for 'Know Your Rights' Notice." Short (300–500 words), deadline-driven, opens with the action item in the first sentence, uses a bulleted list of requirements, links to official agency resources, and closes with a named attorney's direct phone and email.
- **Practice Insight**: e.g., "Privacy Compliance for Growing Companies: What Founders and Employers Need to Know in 2026." Longer (1,200–1,800 words), aimed at business decision-makers, organized under bolded section headings, ends with a "Practical Takeaways" discussion and a "Next Steps" section followed by the standard firm description and office phone list.
- **Doctrinal Deep Dive**: e.g., "Property Searches in Land Use Law Under the Fourth Amendment." Educational walk-through of a legal doctrine, quotes the controlling constitutional or statutory text in italics up front, then traces the case law with full California-style citations (e.g., Camara v. Municipal Court of San Francisco (1967) 387 U.S. 523), and ends with "Next Steps" plus firm boilerplate.

Style fingerprint common to all three

- **Audience:** business owners, employers, founders, and property owners: sophisticated non-lawyers. Plain English with terms of art defined on first use, e.g., "the California Consumer Privacy Act (CCPA)."
- **Voice:** third person and formal but readable; no first-person "I"; the firm appears as "Berliner Cohen" or "our attorneys" only in the closing sections.
- **Structure:** short opening paragraphs that frame why the topic matters now; bolded headings phrased as topics or questions ("Who Is Covered," "What Compliance Actually Requires"); paragraphs of 3–6 sentences; bullets reserved for lists of requirements.
- **California-first:** the analysis centers on California law and its effect on multi-state employers and businesses, often contrasting state law against the fragmented body of federal law.
- **Closers:** a "Next Steps" heading inviting readers to contact the firm; for alerts, a named attorney with phone/email; for longer pieces, the standard paragraph describing the firm, its offices (San Jose, Modesto, Merced), and office phone numbers.
- **Tone of advice:** practical and risk-focused ("far less costly than responding to enforcement actions after the fact") without making promises or giving case-specific legal advice.

Part 2 — Reusable Prompt Components

Build your prompt in two pieces: a style block you reuse every time, plus a format-specific template. Copy the style block below into the start of any article prompt.

The Style Block (paste this every time)

```
You are drafting an article for the Legal News & Insights page of Berliner Cohen
LLP, a Northern California business law firm with offices in San Jose, Modesto,
Merced, and Mariposa.

Write in formal but plain English for business owners and executives, who are
sophisticated non-lawyers. Use third person. Define legal terms and statutes on
first use with parenthetical abbreviations, e.g., "the California Consumer
Privacy Act (CCPA)."

Organize the body under short bolded headings phrased as topics or questions.
Keep paragraphs to 3–6 sentences. Use bullet lists only for enumerable
requirements or deadlines.

Center the analysis on California law and its practical effect on businesses,
noting federal or multi-state context where relevant.

Be practical and risk-focused, but do not promise outcomes, give case-specific
advice, or invent facts, statutes, deadlines, or cases. If you are not certain a
citation or deadline is accurate, flag it with [VERIFY] instead of guessing.

End with a heading titled "Next Steps" inviting readers to contact the firm's
attorneys in the relevant practice area.
```

Template A — Client Alert (300–500 words)

```
[PASTE STYLE BLOCK ABOVE, THEN:]

Draft a client alert titled "California Employer Alert: [DEADLINE/EVENT]." Topic:
[NEW LAW OR RULE, BILL NUMBER, EFFECTIVE DATE].

Open with the action item and deadline in the first sentence. Then add: (1) a
short section explaining the new requirement and which employers it covers; (2) a
bolded heading followed by a bullet list of exactly what must be included or
done; (3) any additional upcoming deadlines under the same law.

Close with: "If your company needs assistance regarding compliance with [LAW],
please contact [ATTORNEY NAME] at [PHONE] or [EMAIL]."

Source material to rely on (do not go beyond it): [PASTE STATUTE TEXT, AGENCY
GUIDANCE, OR YOUR NOTES].
```

Template B — Practice Insight (1,200–1,800 words)

```
[PASTE STYLE BLOCK ABOVE, THEN:]

Draft a practice insight article titled "[TOPIC]: What [AUDIENCE] Need to Know in
[YEAR]."

Open with 2–3 paragraphs framing why this issue now reaches beyond day-to-day
operations (e.g., it surfaces in financing, M&A diligence, enforcement, or
litigation).

Body sections, each under a bolded heading: the state of the law and California's
role; who is covered; what compliance actually requires in operational terms;
enforcement and financial exposure; and any key definitional distinctions readers
commonly get wrong.

Finish with a "Practical Takeaways" section of 2–3 paragraphs of plain-language
guidance, then a "Next Steps" section noting that Berliner Cohen's [PRACTICE
GROUP] attorneys advise clients on these issues.

Key points I want covered: [YOUR 5–8 BULLET NOTES]. Authorities to work from:
[STATUTES/REGS/CASES].
```

Template C — Doctrinal Deep Dive (1,000–1,500 words)

```
[PASTE STYLE BLOCK ABOVE, THEN:]

Draft an educational article titled "[DOCTRINE] Under [SOURCE OF LAW]."

Open by noting that clients commonly ask about this issue, then quote the
controlling constitutional, statutory, or regulatory text in italics.

Walk through the leading authorities chronologically, using California citation
style, in the form Case Name (Year) Volume Reporter Page, with pin cites for
direct quotations. For each case, give the facts in 1–2 sentences, the holding,
and why it matters to property owners or businesses today.

Use ONLY the cases and authorities I provide below; do not add others. Where my
notes are thin, write [EXPAND] rather than inventing detail.

End with a "Next Steps" section on how the firm's attorneys can assist.

Authorities and my case notes: [PASTE YOUR RESEARCH].
```

Part 3 — Your Turn (Exercise)

1. **Pick a topic** from your practice area with a current hook: a new statute, a deadline, a recent decision, or a question clients keep asking.
2. **Choose the format** (A, B, or C) that fits the topic, and gather your source material first: statute text, agency guidance, or your own case notes. The single biggest quality lever is giving the AI your sources rather than letting it rely on memory.
3. **Assemble the prompt** (style block + template) and fill in every bracketed field below before you run it.
4. **Run it, then iterate.** Ask follow-ups like "tighten the opening to one paragraph," "convert the coverage thresholds into a bullet list," or "rewrite the takeaways for hospitality clients."
5. **Review against the checklist** in Part 4 before anything leaves your desk.

Prompt planning fields

Working title: ____________________

Format (A / B / C): ____________________

Audience (e.g., California employers, founders, property owners):

____

Current hook (deadline, new law, recent case):

____

Authorities I will paste in: ____________________

Practice group + contact attorney for the closer:

____

Key points the draft must cover:

____

Part 4 — Attorney Review Checklist

AI output is a first draft, not work product. Before publication:

☐ Every citation verified against the original source: case names, years, volume/page numbers, statute sections, and bill numbers. AI tools can fabricate plausible-looking citations.

☐ Every date and deadline confirmed against the statute or agency notice, not the draft.

☐ All [VERIFY] and [EXPAND] flags resolved.

☐ No statements that could be read as case-specific legal advice or a guarantee of outcome.

☐ Quotes from statutes or opinions checked word-for-word; everything else paraphrased.

☐ Tone matches the samples: plain English, third person, terms defined on first use.

☐ "Next Steps" closer and contact information are correct for the practice group.

☐ No confidential client information was entered into the AI tool, and the tool used is approved under firm policy.

☐ A second attorney in the practice group has reviewed substantive statements of law.

A note on confidentiality and candor: treat anything typed into an AI tool as leaving the building unless the firm's deployment says otherwise, and remember that several courts have sanctioned lawyers for filing unverified AI citations. The same discipline applies to marketing content: an invented case in a client alert is a credibility problem even if it never reaches a docket.

Make Your Own Marketing: Images and Flyers

AI-generated images and one-page flyers for law firm attorneys

Firm marketing runs at two speeds: wait for the designer, or settle for the template everyone else uses. This worksheet adds a third option. You will learn to generate images for articles, posts, and invitations, and to have AI build one-page flyers in HTML or Word.

Two expectations before you start. Flyers come together fast; a usable draft usually appears on the first try and gets good within three rounds. Images take longer. Plan on three to six rounds before one looks right, and treat the rounds as part of the work rather than a sign it isn't working. Budget about 45 minutes for the full worksheet. If you have your firm's logo file and one or two past marketing pieces on your laptop, keep them handy; Part 3 puts them to work.

Part 1 — Which Tool Does What

ChatGPT and Gemini generate images. Claude does not, but it is the strongest of the three at writing the HTML for a designed flyer and at producing a clean Word file. All three can build flyers; only two can make the picture that goes inside one. You may end up using two tools for one piece, and that is normal.

Format	Best for	The catch
AI-generated image	Article headers, LinkedIn posts, the picture inside a flyer or invitation	Takes the most rounds, and text inside images comes out misspelled. Keep the words out of the picture.
HTML flyer	Event flyers and anything that should look designed. The AI controls layout, type, and color completely, so this is where it does its best work.	Editing it next quarter means asking the AI again. You share it by printing to PDF from your browser.
Word flyer	Anything a person at the firm will update by hand: the monthly mixer, the recurring CLE lunch.	Word constrains the design. Expect something cleaner but plainer.

Rule of thumb: if someone will open the file and change the date every month, build it in Word. If it is a one-off that should look sharp, build it in HTML and ship the PDF. Either way, the image inside it comes from Part 3.

Part 2 — The Advertising and Ethics Gate

Marketing changes the rules you have to think about. A flyer is attorney advertising, and AI is happy to decorate one with things that are not true. Clear every box before anything leaves the building.

☐ I am using a firm-approved AI account (a business or enterprise tier where inputs are not used to train models), not a personal free account.

☐ I treat every piece as attorney advertising under California Rules of Professional Conduct 7.1 through 7.3: nothing false or misleading, no guarantees or predictions of results, and the firm identified where the rules require it.

☐ The partner responsible for marketing, or my designated reviewer, approves the final version before it is printed, posted, or mailed.

☐ No client names, matters, logos, or testimonials appear without written consent.

☐ My prompts never ask for a living artist's style, a real person's likeness, or anyone's trademarked characters or trade dress. The firm's own logo is the only logo in the piece.

☐ I understand that AI-generated images may not be copyrightable under current U.S. Copyright Office guidance, so I am not counting on exclusive ownership of the image.

☐ I strip anything the AI added that I did not supply: awards, rankings, “Top 100” badges, statistics, years of experience. The same engine that fabricates plausible-looking case citations, the kind courts have sanctioned lawyers for filing unverified, will fabricate a plausible-looking accolade for your flyer.

Part 3 — Image Generation, Two Tasks

Run these in ChatGPT or Gemini. Two ground rules carry through everything here. First, keep words out of the image; image models misspell with total confidence, and the flyer layer is where text belongs anyway. Second, give the model one concrete visual idea per prompt. "Something professional about estate planning" produces mush. "A single banker's box under a desk lamp" produces a picture.

First, teach it your firm's look

Before generating anything, upload your firm's logo, one or two past marketing pieces, and a screenshot of your website's homepage. Then run this:

```
I've uploaded my law firm's logo, past marketing pieces, and a screenshot of our website.

Write a reusable style guide describing our visual identity: primary and secondary colors as hex codes, typography style (serif or sans, traditional or modern), photography versus illustration, and the overall mood in a few adjectives.

Keep it under 150 words so I can paste it into future image prompts.

Describe only what you can see in the uploads. Do not invent brand elements we don't have.
```

Save the answer somewhere you can reach it. That paragraph is your **style block**, and every prompt below has a slot for it. If you didn't bring brand files today, skip this and fill the slot with three adjectives and a color; you can do it properly back at the office.

Task 1 — A header image for an article or post

Wide format, for the top of a firm article or a LinkedIn post. The concept slot is the one that matters; fill it with something you could photograph.

```
Generate a wide header image, 16:9, for an article titled [ARTICLE TITLE] published by [FIRM NAME], a law firm.

Concept: [ONE CONCRETE VISUAL IDEA, E.G., A SINGLE BANKER'S BOX UNDER A DESK LAMP, OR AN AERIAL VIEW OF A HALF-BUILT SUBDIVISION AT DAWN].

Style: [PASTE YOUR STYLE BLOCK].

Use these colors: [HEX CODES].

No text, letters, or numbers anywhere in the image. I will add the words myself.

No gavels, no scales of justice, no handshakes, no courthouse columns. Restrained and professional, not stock photography.

Do not add any logo, badge, award, or watermark.
```

When a round misses, name the miss. Useful follow-ups:

- “Keep the composition exactly as it is; change the palette to [HEX CODES].”
- “Too busy. Remove the [ELEMENT] and leave more empty space on the left; the title will go there.”
- “Same concept, three variations, different lighting in each.”
- “There is text in the image. Remove every letter and number.”

Task 2 — Imagery for an invitation or flyer

Portrait or square, built to anchor the flyer you will assemble in Part 4.

```
Generate a portrait image, 4:5, to anchor an invitation for [EVENT NAME], a
[EVENT TYPE, E.G., CLIENT RECEPTION, CLE DINNER, WINE TASTING] hosted by [FIRM
NAME], a law firm.

Mood: [E.G., WARM, CANDLELIT, UNHURRIED].

Concept: [E.G., TWO GLASSES OF RED WINE ON A MARBLE COUNTER BESIDE A CLOSED
LAPTOP].

Style: [PASTE YOUR STYLE BLOCK].

No text, letters, or numbers in the image.

No logos, badges, or awards. No recognizable real people.
```

How many rounds is normal: three to six. If round six still isn't close, the problem is the concept, and no pile of adjectives will rescue it. Change the concept and start over. When a round finally lands, ask for it at full resolution and download the file; don't screenshot the preview.

Part 4 — Flyers, Three Tasks

An HTML flyer is a one-page web file the AI designs end to end: layout, type, color, all of it. You preview it in the chat (Claude and ChatGPT both render HTML live), then print it to PDF from your browser with Ctrl+P or Cmd+P. The PDF is what you email, post, or hand to the printer.

Task 1 — The event flyer (HTML)

Every detail is a bracket. Fill them with your next event: the section mixer, the client reception, or a Wine & AI Night of your own.

```
Build a one-page event flyer as a single self-contained HTML file, sized to print on US Letter (8.5 x 11 inches) with everything visible on one page.

Event: [EVENT NAME].

Host: [FIRM NAME].

Date and time: [DATE], [START TIME] to [END TIME].

Location: [VENUE NAME AND ADDRESS].

RSVP: [NAME, EMAIL, AND DEADLINE].

Also include: [E.G., ONE LINE ON WHAT GUESTS WILL DO, PARKING NOTE, CLE CREDIT IF OFFERED].

Design direction: [E.G., ELEGANT AND MINIMAL, DEEP BURGUNDY AND CREAM, SERIF HEADLINES]. Style: [PASTE YOUR STYLE BLOCK IF YOU HAVE ONE].

Leave a clearly marked space near the top for one image; I will supply it.

Use only the details above. If something you need is missing, put a bracketed placeholder in the flyer and tell me what's missing. Never invent a detail, name, or address.

No clip art, no gavels, no fake sponsor logos.
```

Dropping in the Part 3 image: attach the image file to the chat and ask the AI to place it in the space it left. If the platform can't embed your upload directly into the HTML, have it keep the marked box at the image's exact size, then add the picture in your PDF editor, or switch to the Word version below, where inserting a picture is one click.

Iteration lines that earn their keep:

- “Make the event name twice as large and move the RSVP block to the bottom edge.”
- “Swap the palette to [HEX CODES] and lighten the background.”
- “Give me two more directions on the same content: one minimal with lots of white space, one bold and dark.”
- “This won't survive black-and-white printing. Increase the contrast.”

Task 2 — The practice-group flyer (HTML)

One page that answers three client questions: what does this group do, what work does it handle, and who do I call. You supply every fact; the AI supplies the design.

```
Build a one-page flyer as a single self-contained HTML file, US Letter, introducing the [PRACTICE GROUP NAME] group at [FIRM NAME].

Audience: [E.G., GENERAL COUNSEL AT MID-SIZE MANUFACTURERS, REFERRING ATTORNEYS, EXISTING CLIENTS].

Content, in this order: a one-sentence statement of what we do; three to five representative matter types, which I will provide; the attorneys to contact, with the names and titles I provide; office locations.

Here is the raw content: [PASTE YOUR DESCRIPTION, MATTER TYPES, AND CONTACT NAMES].

Design direction: [E.G., CLEAN AND CORPORATE, FIRM COLORS: HEX CODES].

Use only the content I provided. Never add awards, rankings, win rates, statistics, client names, or credentials I didn't give you. If a section feels thin, ask me for more material; don't pad it.

Tighten my wording where it helps, but show me exactly what you changed.
```

Task 3 — The same flyer in Word

Choose Word when a person will maintain the piece: the standing mixer invitation, the CLE series flyer someone re-dates every month. You can build fresh from the Task 1 or Task 2 prompt by asking for "a Word document (.docx) I can download" instead of HTML, or convert what you already built:

```
Here is the HTML flyer we built: [PASTE OR ATTACH IT]. Rebuild it as a .docx file I can edit in Word.

Match the layout, colors, and type as closely as Word allows, and tell me plainly what didn't survive the conversion.

Make the parts I'll edit later (date, time, venue, RSVP) ordinary text, not text boxes buried in graphics.

Use only the details already in the flyer. Anything missing stays bracketed.
```

Platform note: in Claude, switch on file creation in Settings first, then the chat returns a downloadable .docx. ChatGPT produces a downloadable file in the chat. Gemini drafts the content and exports it to Google Docs, which will save as .docx.

Part 5 — Hands-On: Build the Next Flyer

Pick something real: the section mixer, the CLE lunch, the holiday party, or your group's one-pager. A real piece with a real date keeps you honest about the details, and you leave with work product instead of practice.

The piece you're building: ____

The details (date, time, venue, RSVP, contacts):

____

Format and why (HTML, Word, or image only):

____

Brand assets to upload (logo, past pieces, site screenshot):

____

Design direction in five words or fewer:

____

Now run it:

1. Clear the gate in Part 2. Every box.
2. Upload your brand assets and run the style-guide prompt from Part 3. Save the style block.
3. Optional: generate the image (Task 1 or 2). Cap yourself at six rounds today.
4. Paste the matching flyer prompt from Part 4 with your details filled in.
5. Iterate at least three rounds. Name what's wrong precisely each time; “make it pop” gets you nothing.
6. Produce the final file: print the HTML to PDF, or download the .docx.
7. Read every word aloud. Check the date against your calendar, the address against a map, and the RSVP address by sending it a test email.
8. Save the work. Ask: “Write me a reusable prompt that produces this flyer format with new event details.” Then create a project (Claude or ChatGPT) or a Gem (Gemini) named “Firm Marketing,” paste that prompt into its instructions, and upload your style block and logo as its knowledge. The next event starts already set up.

Before Anything Leaves the Building

☐ A human read every rendered word, including any text inside images. Image models misspell with total confidence.

☐ The reviewer responsible for firm advertising approved the piece under Rules of Professional Conduct 7.1 through 7.3.

☐ Date, time, address, and RSVP details were verified against the calendar invite, the venue's website, and a test email.

☐ Nothing in the piece was invented by the AI: no surprise awards, rankings, statistics, or credentials.

☐ No client confidences, names, or matters appear without written consent.

☐ You, not the AI and not a designer, signed off as the final reviewer.

The payoff: a finished one-page flyer is roughly twenty minutes of work, most of it spent choosing between two palettes the AI offered. The second flyer is faster than the first, because by then the style block, the prompt, and the project already exist, and the AI starts out knowing what your firm looks like.

Build Your Own Law Firm Website with Claude

A Step-by-Step Worksheet for Attorneys (No Tech Experience Required)

From blank page to live website in one afternoon

This worksheet walks you, step by step, through building a professional website for your practice and putting it on the internet where clients can find it. You will use Claude (an AI assistant) to design and write the website, GitHub (a free service) to host it, and a domain registrar to buy your web address (like www.smithlawfirm.com).

Total cost: roughly $10–$25 per year for your domain name. The website design and hosting are free.

Total time: about 2–3 hours, and you do not need to finish in one sitting. Check off each step as you complete it.

THE MOST IMPORTANT INSTRUCTION IN THIS WORKSHEET

You are never stuck. If anything in this guide confuses you, if a screen does not look the way this worksheet says it should, or if something just is not working:

1. Open Claude (claude.ai) or your AI assistant of choice.
2. Upload this worksheet to the chat (click the paperclip / plus icon and attach this file).
3. Type: *"I'm using this worksheet to build my website. I'm on Step ___ and here's what's happening..."* and describe the problem. You can also paste a screenshot directly into the chat.

The AI will read the worksheet, see exactly where you are, and guide you through it in plain English. Use this trick as often as you need — that is what it is for.

Before You Begin: What You Need

- ☐ A computer with an internet browser (Chrome, Safari, Edge, or Firefox)
- ☐ An email address you check regularly
- ☐ A credit card (only for the domain name, about $10–$25/year)
- ☐ A Claude account — sign up free at claude.ai (a paid plan helps for longer sessions but is not required)
- ☐ Optional but recommended: a professional headshot and/or firm logo saved on your computer

A quick note on attorney advertising rules: in most jurisdictions, a law firm website is attorney advertising. Before you publish, check your state bar's rules on required disclaimers (e.g., "Attorney Advertising," "Prior results do not guarantee a similar outcome"), claims of specialization, and testimonials. Step 1 includes a place to note what your jurisdiction requires — and you can ask Claude to summarize your state's rules as a starting point for your own verification.

Step 1: Gather Your Content (15–30 minutes)

A good website starts with good raw material. Fill in the blanks below — you will paste this information into Claude in Step 2. Do not worry about polish; Claude will turn your notes into professional website copy.

Firm / practice name: ______________________________

Your name and title: ______________________________

Tagline (optional, e.g., "Trusted family law counsel in San Jose"): ______

Practice areas (list 2–6): __________________________

City / region you serve: ____________________________

Phone number: ___________________________________

Email for client inquiries: __________________________

Office address (if you want it listed): ________________

Bar admissions (states, year admitted): ________________

Education and notable experience: ____________________

Required disclaimers per your state bar: _______________

Also jot a few sentences about: who your ideal client is, what makes your practice different, and how you want the site to feel (e.g., "modern and approachable," "traditional and authoritative," "warm and personal"). Note any colors you like.

> **IF YOU GET STUCK:** Not sure what to write? Ask Claude: "I'm a [practice area] attorney in [city]. Interview me one question at a time to gather everything needed for a great law firm website bio and homepage." It will ask questions and draft the content with you.

Step 2: Have Claude Build Your Website (30–60 minutes)

1. Go to **claude.ai** in your browser and sign in (or create a free account).
2. Start a new chat and paste the prompt below, replacing the bracketed parts with your notes from Step 1.

COPY AND PASTE THIS PROMPT INTO CLAUDE:

I am an attorney with no technical experience, and I want you to build me a complete, professional one-page website for my law practice. Here are my details:

[PASTE YOUR STEP 1 NOTES HERE — firm name, your bio, practice areas, contact info, bar admissions, required disclaimers, and the style/feel you want]

Please build it as a single, self-contained HTML file with all styling included in the file, so it works on its own with no other files needed. It should look clean and professional, work well on phones as well as computers, and include: a header with my firm name, sections for my practice areas, an about/bio section, a contact section with my phone and email, and a footer with my required attorney advertising disclaimer. Show me the result so I can see it, and then I will ask for changes.

3. Claude will create the website and display a live preview next to the chat (this preview panel is called an "artifact"). Look it over like a draft brief.
4. Ask for revisions in plain English, one or two at a time. For example: *"Make the colors navy and gold," "Make my name bigger," "Add a section about free consultations," "Make it feel more modern."* Repeat until you are happy.
5. If you have a headshot or logo, you can attach the image and ask Claude where it could go — or simply skip photos for version one and add them later. A clean text-only site is perfectly professional.
6. When you are satisfied, tell Claude: *"This is final. Please make sure it is one single HTML file and name it index.html."*
7. Download the file: at the top-right of the preview panel, click the download icon (or the menu next to it) and save the file to your computer. **It must be named exactly index.html** — all lowercase. If your computer saved it under a different name, rename it. Remember which folder it is in (usually Downloads).

☐ I downloaded a file named index.html to my computer

IF YOU GET STUCK: Can't find the download button, or your file has the wrong name? Upload this worksheet to Claude and say: "I'm on Step 2, item 7 — help me download my website as index.html."

Step 3: Create a Free GitHub Account (10 minutes)

GitHub is a website that stores files and — importantly for you — will host your website on the internet for free through a feature called GitHub Pages. Lawyers, businesses, and millions of professionals use it. You only need the basics.

1. Go to **github.com** and click **Sign up**.
2. Enter your email, create a password, and choose a username. **Choose your username carefully** — it becomes part of your free website address. Something professional like "smithlawoffice" works well (lowercase letters, numbers, and hyphens only).
3. Verify your email by clicking the link GitHub sends you, and complete any puzzle it asks you to solve.
4. Choose the **Free** plan if asked. You never need a paid plan for this.

My GitHub username: ______________________________

☐ I created my GitHub account and verified my email

> **IF YOU GET STUCK:** Stuck on this step? Open Claude (or your AI of choice), upload this worksheet, and type: "I'm working through this worksheet and I'm stuck on this step. Here's what I see on my screen..." Then describe (or screenshot) what's happening. The AI will walk you through it.

Step 4: Put Your Website File on GitHub (15 minutes)

Now you will create a "repository" — think of it as a folder on GitHub — and upload your index.html file into it. Everything is done in the browser by pointing and clicking. No coding, no command line.

1. While signed in to GitHub, click the **+** icon in the top-right corner and choose **New repository**.
2. In the "Repository name" box, type **your-username.github.io** — replacing "your-username" with your exact GitHub username from Step 3. Example: if your username is smithlawoffice, type *smithlawoffice.github.io*. This special name is what makes GitHub treat it as a website.
3. Make sure **Public** is selected (free website hosting requires a public repository — only your website file is public, nothing else).
4. Click the green **Create repository** button.
5. On the next page, click the link that says **uploading an existing file** (it is in the small text near the middle of the page).

6. Drag your **index.html** file from your Downloads folder onto the page (or click "choose your files" and select it).
7. Click the green **Commit changes** button. ("Commit" is just GitHub's word for "save.")

☐ My repository is named exactly my-username.github.io

☐ My index.html file appears in the repository's file list

> **IF YOU GET STUCK:** Stuck on this step? Open Claude (or your AI of choice), upload this worksheet, and type: "I'm working through this worksheet and I'm stuck on this step. Here's what I see on my screen…" Then describe (or screenshot) what's happening. The AI will walk you through it.

Step 5: See Your Website Live on the Internet (5–10 minutes)

Because you used the special repository name, GitHub automatically publishes your site. Let's confirm.

1. Wait 2–5 minutes after uploading. GitHub needs a moment to publish.
2. Open a new browser tab and go to: **https://your-username.github.io** (again, replacing "your-username"). Your website should appear!
3. If you see a "404" error, wait five more minutes and refresh. Still nothing? In your repository, click **Settings** (top menu), then **Pages** (left sidebar). Under "Branch," make sure it says **main** and click Save. Check again in a few minutes.

☐ I can see my website at https://my-username.github.io

Milestone! Your website is now live on the internet. Anyone in the world can visit it at that address. The remaining steps simply give it a more professional address of your own, like www.smithlawfirm.com.

> **IF YOU GET STUCK:** Stuck on this step? Open Claude (or your AI of choice), upload this worksheet, and type: "I'm working through this worksheet and I'm stuck on this step. Here's what I see on my screen…" Then describe (or screenshot) what's happening. The AI will walk you through it.

Step 6: Buy Your Own Domain Name (15 minutes)

A domain is your web address. You rent it annually from a company called a registrar. Well-known, reasonably priced registrars include Namecheap, Cloudflare Registrar, Porkbun, and GoDaddy — they all work for this guide. Expect roughly $10–$25 per year for a .com.

1. Decide on a name. Tips: shorter is better; .com is the safest choice (though .law and .legal exist at higher prices); avoid hyphens and hard-to-spell words; consider yourname + law, e.g., smithlaw.com or smithfamilylaw.com.

Domain names I'd like, in order of preference (1): ____________

(2): ____

(3): ____

2. Go to a registrar's website (e.g., **namecheap.com**), type your desired name in the search box, and see if it is available. If taken, try your next choice.
3. Add the domain to your cart and check out. **Decline the add-ons** (hosting, email, website builders, "premium DNS") — you do not need them. Free WHOIS/domain privacy is worth keeping if offered.
4. Create an account with the registrar during checkout and save your login details somewhere safe — you will need them in Step 7.

The domain I purchased: ________________________________

Registrar I bought it from / login email: __________________

☐ I own my domain and can sign in to my registrar account

> **IF YOU GET STUCK:** Registrar websites love to upsell. If you're unsure whether you need something at checkout, ask your AI: "I'm buying a domain at [registrar] for a website hosted on GitHub Pages. Do I need [name of add-on]?" The answer is almost always no.

Step 7: Connect Your Domain to Your Website (20–30 minutes + waiting)

This is the most technical-feeling step, but it is really just copying five short entries into a form. You will tell GitHub about your domain, then tell your domain where to find GitHub. Take it slowly.

Part A — Tell GitHub your domain

1. On GitHub, open your repository, click **Settings**, then **Pages** in the left sidebar.

2. Find the box labeled **Custom domain**, type your domain with www in front (e.g., *www.smithlaw.com*), and click **Save**. A "DNS check" message may show an error at first — that is expected until Part B is done.

Part B — Tell your domain where GitHub is (DNS records)

3. Sign in to your registrar (Namecheap, etc.), find your domain, and look for a section called **DNS**, **DNS settings**, **Advanced DNS**, or **Manage DNS records**.
4. If there are existing records of type "A" or "CNAME" pointing to the registrar's "parking page," delete those. Leave any records of other types (like NS or SOA) alone.
5. Add the five records in the table below, one at a time. "Host" is sometimes labeled "Name." The @ symbol means "the domain itself." If asked for TTL, leave the default.

Type	Host / Name	Value / Points to
A	@	185.199.108.153
A	@	185.199.109.153
A	@	185.199.110.153
A	@	185.199.111.153
CNAME	www	your-username.github.io (your actual GitHub username)

6. Save the records, then wait. DNS changes usually take 15–60 minutes to take effect, occasionally up to 24 hours. This is normal — a perfect time for a coffee or a billable hour.

Part C — Turn on the security padlock (HTTPS)

7. After the wait, return to GitHub → your repository → Settings → Pages. The DNS check next to your custom domain should now show a green success message.
8. Check the box labeled **Enforce HTTPS**. (If it is grayed out, GitHub is still issuing your free security certificate — check back in an hour.) This gives visitors the padlock icon and keeps browsers from flagging your site.
9. Type your domain (www.yourdomain.com) into a browser. Your website should appear at your own professional address.

☐ All five DNS records are saved at my registrar

☐ GitHub shows a green DNS check and "Enforce HTTPS" is checked

☐ My website loads at my own domain

IF YOU GET STUCK: DNS is the step where most people want backup. Upload this worksheet to your AI and say: “I’m on Step 7 using [registrar name]. Walk me through adding these exact DNS records on their website.” Screenshots of your registrar’s DNS page help enormously.

Step 8: Add a Contact Form — Get Notified When Someone Needs an Attorney (30–45 minutes)

Right now, visitors can call or email you. A contact form is even better: a potential client fills in their name and a short description of their issue, clicks Submit, and you get an email alert immediately — a hot lead delivered to your inbox.

One wrinkle: GitHub Pages hosts simple websites and cannot process form submissions by itself. The standard solution is a free service called Formspree, which receives the form and emails it to you. Its free plan includes 50 submissions per month — plenty for most practices — and you only sign up and copy one web address. Claude does the rest.

Part A — Set up Formspree (your form's mailbox)

1. Go to **formspree.io** and click **Sign up**. Use the email address where you want client inquiries delivered, and verify it when Formspree sends you a confirmation email.
2. Once signed in, create a new form (look for a button like **+ New form**). Name it something like "Website Inquiries" and confirm the email address that should receive submissions.
3. Formspree will show you your form's **endpoint** — a web address that looks like *https://formspree.io/f/abcdwxyz*. Copy it. This is the only technical ingredient you need.

My Formspree endpoint address: ______________________________

Part B — Have Claude add the form to your website

4. Start a new Claude chat and attach your current **index.html** file (download it from your GitHub repository first if you don't have a copy: open the repository, click index.html, then use the download button near the top-right of the file view).
5. Paste the prompt below, inserting your Formspree address where indicated.

COPY AND PASTE THIS PROMPT INTO CLAUDE:

This is my live law firm website. Please add a professional "Contact Us" section with a form that submits to this Formspree endpoint: [PASTE YOUR FORMSPREE ADDRESS HERE]

The form should ask for: full name, email, phone number, and a brief description of their legal matter. After someone submits, they should see a friendly confirmation message on the page (do not send them away to another site).

Important: directly above the form, include this notice in smaller text: "Submitting this form does not create an attorney-client relationship. Please do not include confidential or time-sensitive information."

Keep everything else about the site exactly the same, match the form's style to the rest of the design, and give me back the complete updated file named index.html.

6. Review the preview, ask for any tweaks, then download the updated **index.html**.
7. Upload it to GitHub just like in Step 4: open your repository, click **Add file → Upload files**, drag the new index.html in, and click **Commit changes**. It replaces the old version, and your live site updates within a couple of minutes.

Part C — Test it

8. Visit your live website, fill out your own form with a test inquiry, and click Submit.
9. First submission only: Formspree may email you a link to confirm or activate the form. Click it, then submit one more test.
10. Check your inbox for the submission email. **Check your spam folder too** — if it landed there, mark it "Not spam" and add the Formspree sender to your contacts so you never miss a real client.
11. Submit a test from your phone as well, since many clients will find you that way.

- ☐ My form appears on the live site and matches the design
- ☐ A test submission arrived in my email inbox (not spam)
- ☐ The no-attorney-client-relationship notice appears above the form

Two practice pointers: First, treat form notifications like voicemail — check that inbox routinely, and remember the form is for initial contact only, not confidential client communications. Second, if your practice grows past 50 inquiries a month (a good problem), Formspree's paid plan raises the limit, or you can ask your AI about alternatives.

> **IF YOU GET STUCK:** Form not working? Upload this worksheet plus your index.html to your AI and say: "I'm on Step 8. My form [doesn't appear / shows an error / submissions aren't arriving by email]." Include your Formspree address and a screenshot if you can.

Step 9: Making Changes Later (and Keeping It Current)

Your website is live — congratulations, counselor. When you want to update it (new practice area, new phone number, a results page), here is the simple loop:

1. Go to your GitHub repository, click on **index.html**, and click the small **download/raw or copy button** — or simply paste your live website address into Claude and say what you want changed.
2. Easiest method: start a new Claude chat, attach your current index.html (or paste its contents), and say: *"Here is my current law firm website. Please [describe the change] and give me back the complete updated index.html file."*
3. Download the updated file (still named index.html), then in your GitHub repository click **Add file → Upload files**, drag the new file in, and click **Commit changes**. The new version replaces the old one, and your live site updates within a couple of minutes.

Annual maintenance — mark your calendar

- **Domain renewal:** your registrar will charge the yearly fee. Turn on auto-renew so your address never lapses (lapsed legal domains get snapped up quickly).
- **Content check:** review bar disclaimers, contact info, and practice areas once or twice a year.

Troubleshooting quick reference

What you see	What it usually means / what to do
404 page at your-username.github.io	Repository name doesn't exactly match username.github.io, or the file isn't named index.html (all lowercase). Fix the name, wait 5 minutes.
"DNS check unsuccessful" on GitHub	Records not saved correctly or still propagating. Re-check the five records against the Step 7 table, wait an hour, then click Save on the custom domain again.
Site works at github.io but not your domain	DNS still propagating (up to 24 hrs) or a leftover old record at the registrar is interfering. Delete conflicting A/CNAME records.
"Not secure" warning in browser	Enforce HTTPS isn't on yet. GitHub → Settings → Pages → check Enforce HTTPS once available.
Website looks wrong on phones	Ask Claude: "Make my site fully mobile-friendly" and re-upload the new index.html.

Form submissions not arriving by email	Check spam folder first. Then sign in to formspree.io — submissions are listed there; confirm the form is active and the notification email is verified.
Anything else	Upload this worksheet to your AI, state the step number, describe or screenshot the problem.

Final checklist

- ☐ Website content reviewed for accuracy and bar-rule compliance
- ☐ Site loads at my custom domain with the padlock (HTTPS)
- ☐ Phone number and email on the site are correct (test them!)
- ☐ Contact form tested — a submission arrived in my inbox
- ☐ Site checked on my own phone
- ☐ Domain auto-renew turned on at my registrar
- ☐ New domain added to my email signature, business cards, and bar/legal directory listings

You did it. You designed, built, and launched a professional website — and you now know how to maintain it yourself for the cost of a single client lunch per year. Whenever you want to expand it (a blog, client testimonials, a contact form, separate pages per practice area), just ask Claude: it built the site, and it can grow it with you.

Build a Client Pitch Deck with Claude Design

Turn a case into a polished, on-brand pitch in one sitting — a litigation walkthrough

A new-client pitch is mostly judgment and a little design. The judgment is yours; the design is what eats your evening. **Claude Design** closes that gap: you describe the engagement in plain language and it builds a complete, on-brand slide deck — strategy, timeline, team, fees, and speaker notes — in minutes. You refine it by talking to it, then export a PowerPoint.

This worksheet walks the whole loop using a real example we built live: a fictional litigation boutique, **Calderon Hale LLP**, pitching to represent **Meridian Healthcare Systems** as plaintiff against its software vendor, Quantum Dynamics — the same matter from the Discovery and Complaint builder worksheets. Every slide pictured below came out of one prompt.

Before You Start: The Confidentiality Gate

A pitch deck is marketing, not a filing, but it still touches real matter facts. Walk the gate before you paste anything in.

> - ☐ **Firm-approved AI account only**, on a plan where your inputs are not used to train models. Never a personal free account.
> - ☐ **Keep privileged strategy out of the prompt.** Pitch at the level a prospective client could already see: the dispute, your approach, your record. Strip case numbers and sensitive specifics.
> - ☐ **Figures and names are placeholders until you say otherwise.** The model will invent reasonable-looking numbers (losses, recoveries, fee ranges). Replace them with real, verified figures before the deck leaves your office.
> - ☐ **Mind the marketing rules.** Track-record claims need a prior-results disclaimer, and California attorneys should check the State Bar's Practical Guidance on generative AI before relying on any of this.

What Claude Design Is

Claude Design (a research preview from Anthropic Labs) lives at **claude.ai/design** and is included with Claude Pro, Max, Team, and Enterprise plans. (On Enterprise it is off until an admin enables it.) You describe what you want; Claude builds a first version; you refine it by conversation, direct edits, or sliders. Decks export to **PowerPoint (.pptx), PDF, standalone HTML, or Canva.**

The Walkthrough

Step 1 — Open a new slide deck

1. Go to **claude.ai/design** and sign in with your firm account.
2. Choose the **Slide deck** tab, give the project a name (e.g., "Meridian Litigation Pitch"), and turn on **Use speaker notes** so every slide comes with talking points.
3. Click **Create**. You land in a workspace with a "Describe what you want to create…" box on the left and a canvas on the right.

Step 2 — Describe the deck

This is the whole game. A good deck prompt names five things: **who is presenting, who is in the room, the matter, the slide list, and the brand.** Paste it as one message (in this composer, Enter sends and Shift+Enter makes a new line). Here is the exact prompt we used — swap the bracketed parts for your own engagement:

> Create a polished 10-slide litigation pitch deck for a law firm pitching to win a new client engagement.
>
> PRESENTING FIRM: [your firm], a commercial litigation boutique.
>
> AUDIENCE (prospective client): [client], its General Counsel and executive team.
>
> THE MATTER: [one plain-English paragraph — the dispute, who did what, and what you would pursue].
>
> SLIDES: (1) Title; (2) We understand your situation; (3) The legal theory; (4) Our litigation strategy, phased; (5) Roadmap & timeline; (6) Your team; (7) Track record; (8) Fee structure — a transparent phased budget with an alternative-fee option, and a line on how AI-driven drafting keeps the budget predictable; (9) How we use AI litigation tools (e.g., Legion) to lower and stabilize cost, research more thoroughly, and free our lawyers for strategy instead of drafting — note that every AI draft and authority is attorney-reviewed; (10) Next steps and contact.
>
> BRAND: refined, confident law-firm aesthetic; primary deep maroon #610338, near-black text, secondary gray, light backgrounds with subtle accents, clean sans-serif, generous whitespace, professional, not flashy.
>
> Add concise speaker notes with talking points for each slide.

Step 3 — Let Claude build and check its own work

Claude writes a short plan, sets up the deck, and builds all ten slides — then it screenshots its own output and fixes what looks wrong. In our run it caught an entrance animation that left text invisible (which would also have broken the PDF export) and removed it on its own. Give it a few minutes; you do not have to babysit it.

Step 4 — Review and refine

You have four ways to change anything, from lightest touch to most precise:

- **Tweaks** — a "Describe a tweak..." box for plain-language nudges ("make the title slide less heavy," "tighten the timeline").
- **Edit** — click any heading or paragraph to retype it directly; the panel also sets page background, font, and base size.
- **Comments** — mark up specific elements and batch your notes, then have Claude apply them at once.
- **Just ask in chat** — "Add a slide on e-discovery scope," "swap slides 6 and 7," "make the fee table an alternative-fee comparison."

Step 5 — Keep it on brand

The fastest route is to name your colors in the prompt, as above. For repeat use, click **Set up design system** once and load your firm's palette and type. After that, every deck anyone on the team builds comes out in your brand automatically.

Step 6 — Export to PowerPoint

1. Ask in chat: *"Export this deck to PowerPoint (.pptx), and also produce a PDF version."*
2. The **.pptx downloads with editable text, shapes, and all your speaker notes attached** — open it in PowerPoint and finish there. For the PDF, a print-ready tab opens; press Ctrl+P and Save as PDF (margins None, background graphics on, so the maroon slides print).
3. You can also right-click the page in **Design Files** and choose **Download** for the standalone HTML version.

What One Prompt Produced

Every slide below is straight from the export — no manual design. This is the Calderon Hale pitch for the Meridian matter.

CALDERON
HALE

COMMERCIAL LITIGATION

Litigation Representation Proposal

Prepared for Meridian Healthcare Systems, LLC

Re: Meridian Healthcare Systems, LLC v. Quantum Dynamics Software, Inc. — breach of contract, breach of express and implied warranties, and fraudulent inducement.

June 2026

Privileged & Confidential — Attorney Work Product

Slide 1 — Title. Firm, client, the matter, and a privileged/work-product footer.

THE SITUATION

A platform at the center of your operations failed in production.

Meridian licensed Quantum's scheduling and EHR platform to run core clinical operations across its facilities. The system failed repeatedly in production — scheduling outages, lost records, and data-integrity failures — disrupting patient care and forcing costly manual workarounds.

Quantum represented an enterprise-ready platform and has not made Meridian whole. The losses are documented, the contract terms are clear, and the window to act decisively is now.

11 mo.
of recurring production failures before the platform was pulled

23
Meridian facilities running on the Quantum platform

$18M+
estimated operational losses and remediation costs

3
strong, independent claims supporting recovery

CALDERON
HALE

02

Slide 2 — "We understand your situation," with stat cards framing the loss.

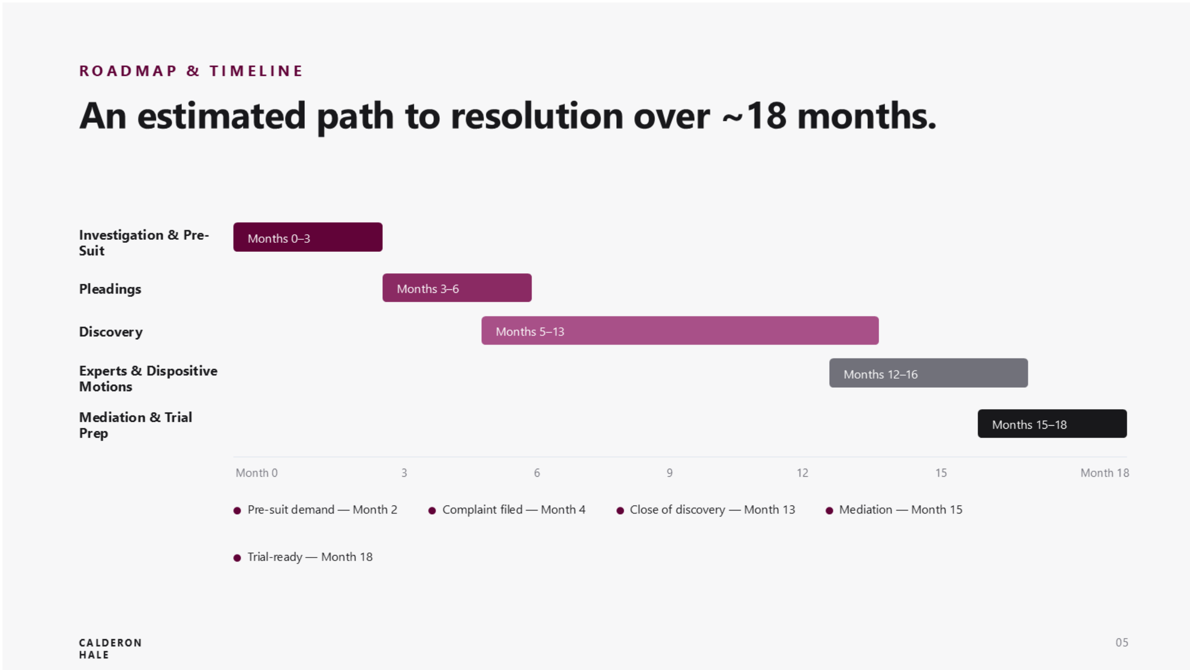

Slide 5 — An 18-month litigation roadmap as a phase-bar timeline with milestones.

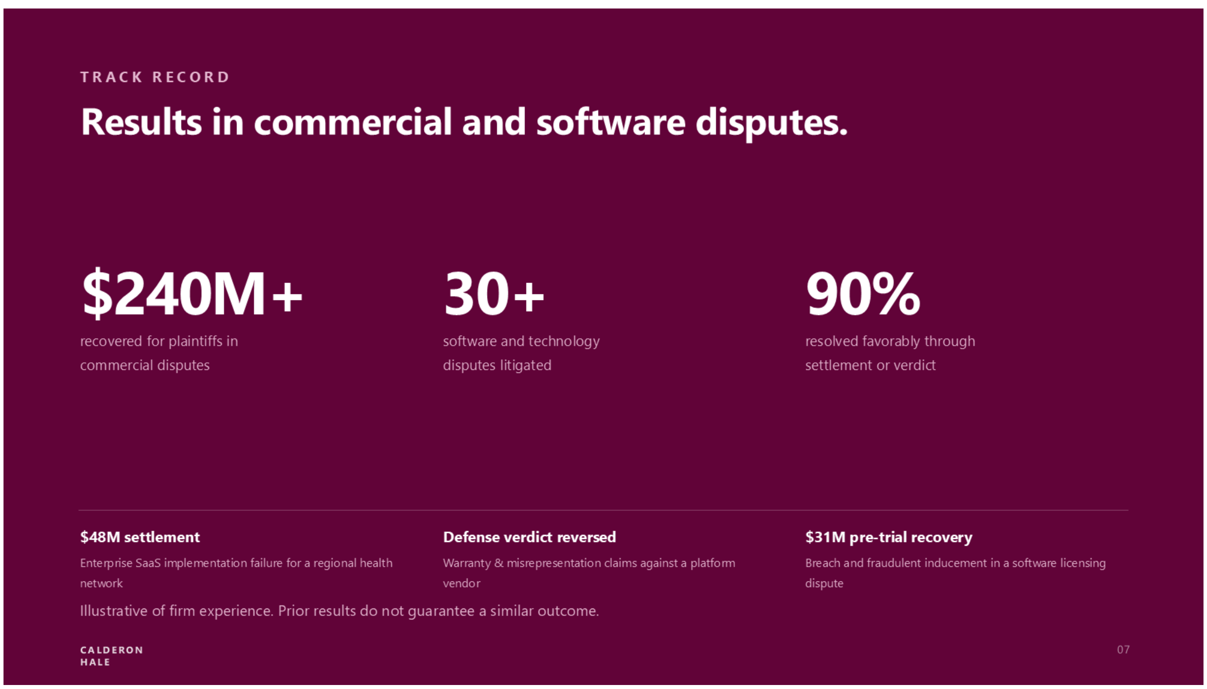

Slide 7 — Track record, including a built-in prior-results disclaimer.

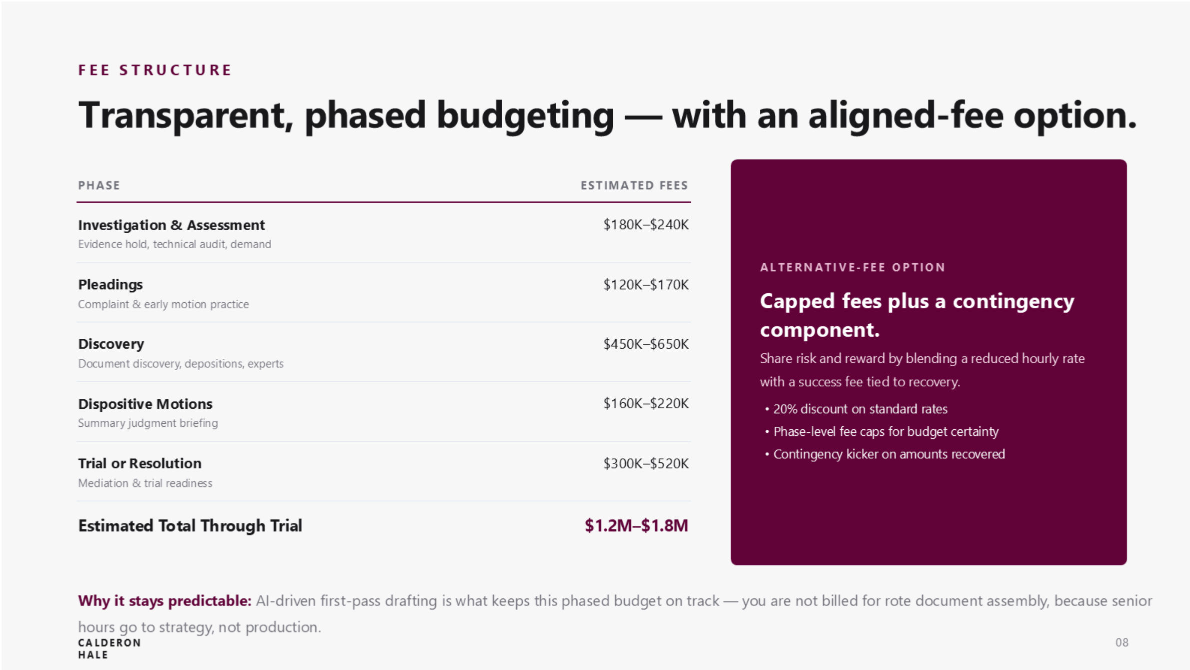

Slide 8 — A phased budget and an alternative-fee option, with a line tying budget predictability to AI-driven drafting.

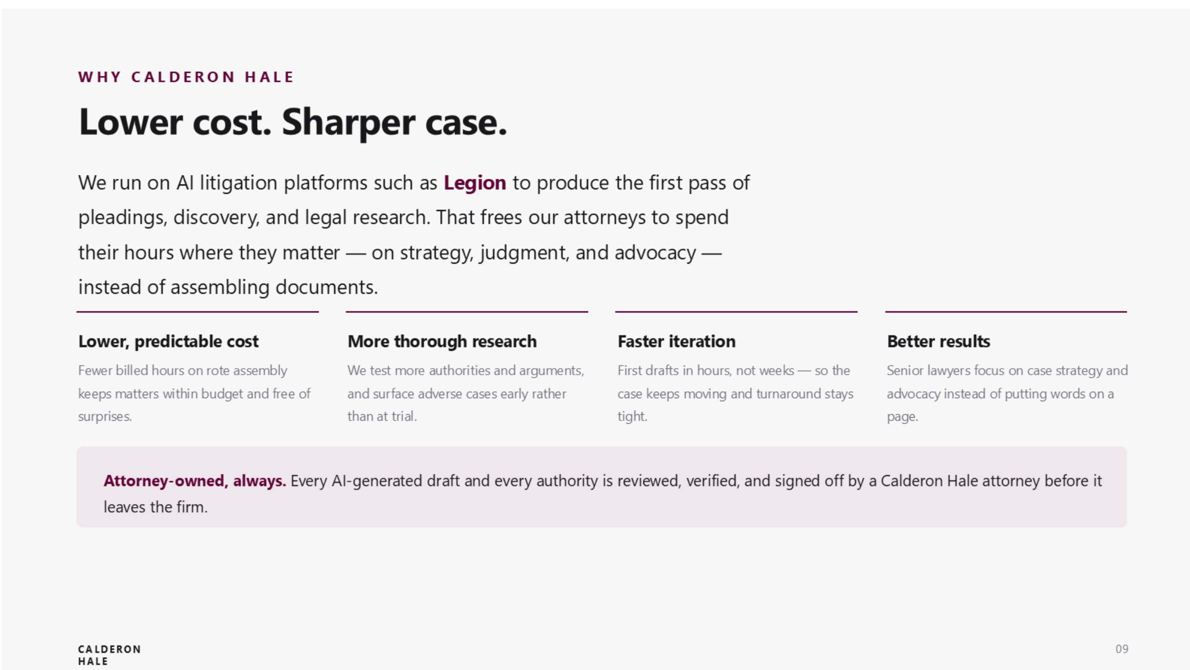

Slide 9 — The cost-and-results story: how the firm uses AI tools like Legion so lawyers spend their hours on strategy, not drafting.

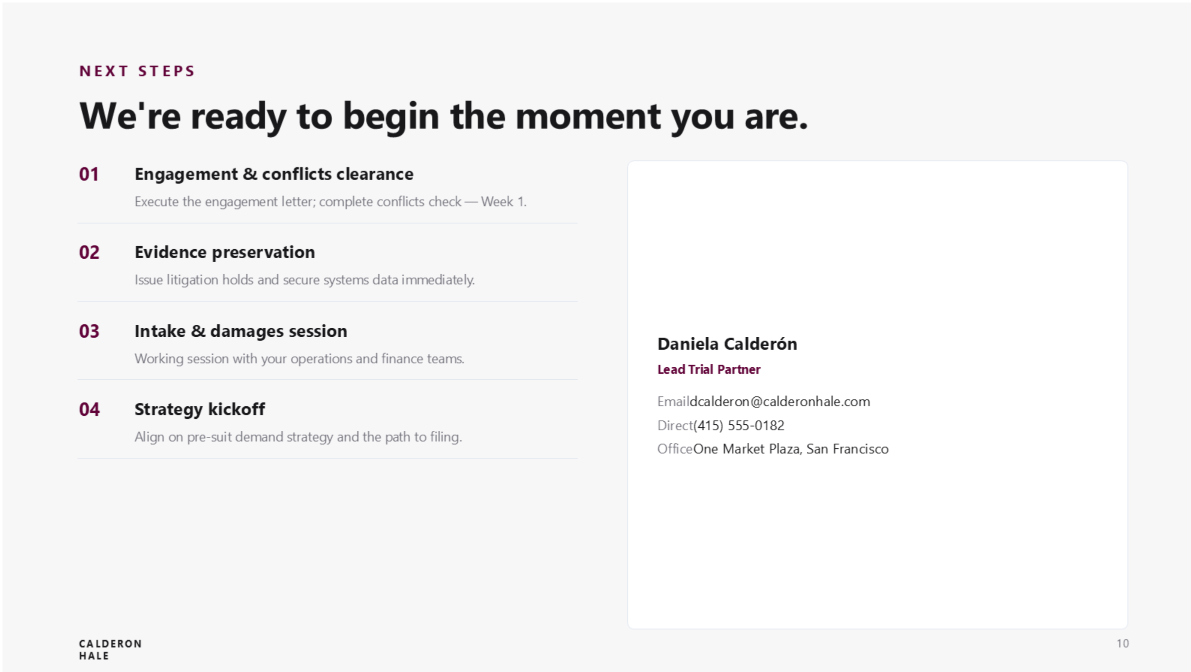

Slide 10 — Next steps and a clean contact card to close the pitch.

Make your AI advantage explicit

General counsel are under real pressure to control outside-counsel spend, so do not bury how you work. A dedicated slide that says you use AI litigation tools to draft the first pass — keeping the budget predictable, the research thorough, and senior lawyers on strategy instead of document assembly — is often the most persuasive slide in the deck. Pair it with a clear line that every draft and authority is attorney-reviewed, so efficiency never reads as corner-cutting.

Your Turn

Pick a real (or realistic) matter you could pitch and build it now. Work through these in order:

1. Name the **presenting firm** and the **audience** (who is actually in the room).
2. Write the **matter paragraph** in plain English — the dispute, the stakes, and what you would pursue.
3. Paste the Step 2 prompt with your details, then let it build.
4. Make two refinements with **Tweaks** and one direct fix with **Edit**.
5. Replace every placeholder figure and name with real, verified numbers.
6. Export to PowerPoint and finish in your own template.

The one habit that matters

Once a deck format works, save the prompt. Next pitch, you change the firm, client, and matter paragraph and you are ninety percent done before you open PowerPoint. That is the whole point: build the frame once, reuse it forever.