ACTIONABLE AI

presents

AI Playbook

Claude Chat and Cowork for Small Law Firms

A Practical Guide to Using Claude Responsibly, Ethically, and Effectively in Your Day-to-Day Practice

actionableai.net · Prepared by Greg Yates - Revised September 29, 2026

Important Notice

This playbook is educational and informational. It is not legal advice, ethics counsel, or a substitute for consultation with The Florida Bar or qualified counsel on your firm's specific professional-responsibility obligations.

Claude and the other products named in these pages are named because small firms already own or evaluate them, not as endorsements. Actionable AI takes no fees from any vendor and recommends no product.

Nothing here creates an attorney-client relationship. Actionable AI is a consulting firm, not a law firm.

Claude's product features, plan terms, pricing, and data-handling policies change. Verify current features and terms of service at anthropic.com and support.claude.com before making tool-selection decisions based on any guidance here. Florida Bar Ethics Opinion 24-1 (January 19, 2024) is the primary Florida authority on attorney AI use. Read it in full at floridabar.org; the summary in this playbook is not a substitute.

All sample governance policies, client disclosures, certification language, and prompts are starting points. Have a Florida-licensed attorney review them before adoption or use in a client matter.

A note on numbers. Legal amounts, deadlines, and case outcomes carry citations; third-party statistics name their source. Other figures are Actionable AI working assumptions. Test those against your firm's rates, volume, and matter values before repeating them as fact.

Contents

Important Notice....1
The Problem With Generic AI Advice....4
The Cost of Wrong....4
What This Playbook Does Differently....5
The Actionable AI Philosophy....5
Why Small Law Firms Need a Different AI Approach....6
The Non-Technical Partner Is the Decision-Maker....6
Three Claude Interfaces, and What This Playbook Covers....6
A Note on Claude Code....7
Claude Chat: The Conversational Interface....7
Claude Cowork: The Agentic Assistant....8
The Confidentiality Imperative: Why Your Plan Choice Matters....8
The Florida Ethics Reality....9
The Five Rules That Matter Most....9
Other Rules That Touch This Work....10
The Courts Are Writing Their Own Rules: Right Now....11
The Emerging Issue Most Small Firms Are Missing: From AI Chat to AI Agent....12
What an AI Agent Actually Does....12
The Scale of the Shift....12
Three Immediate Opportunities....12
The Governance Implication....13
Ethics Risks and Governance Framework....14
High-Risk Failure Modes....14
Data Classification Framework....14
Approved Tool List and Governance Policy....15
Staff Training Requirements....15
Best Use Cases by Workflow Stage....16
Stage 1: Client Intake and Intake Analysis....16
Stage 2: Legal Research....16
Stage 3: Document Drafting....17
Stage 4: Document Review and Quality Check....17
Stage 5: Client Communications....17
Stage 6: Matter Administration....18
Stage 7: Marketing and Business Development....18
Prompts Library....19

Stage 1: Intake and Intake Analysis....19
Stage 2: Legal Research....20
Stage 3: Document Drafting....21
Stage 4: Document Review....22
Stage 5: Client Communications....22
Stage 6: Matter Administration....23
Stage 7: Marketing and Business Development....23
Appendix A: Sample AI Governance Policy....25
1. Purpose....25
2. General Principles....25
3. Definitions....25
4. Permitted Uses....25
5. Prohibited Uses....26
6. Confidentiality and Data Handling....26
7. Claude Cowork: Agent Policy....26
8. Court and Tribunal Filings....26
9. Incident Response....27
10. Policy Acknowledgment....27
Appendix B: Client AI Disclosure Statement....28
Appendix C: Court AI Disclosure and Certification....29
Appendix D: Letter to AI Vendors: Nine Questions....30
Appendix E: 90-Day Implementation Roadmap....31
Days 1-30: Foundation....31
Days 31-60: Proficiency....31
Days 61-90: Optimize....31
Sources and Ethics Anchors....32
About Actionable AI....33
What We Do....33
Who This Playbook Is For....33

The Problem With Generic AI Advice

Here is what the AI influencers will not tell you: the advice flooding your inbox was written for tech companies, startups, and marketing agencies. The author has never pulled an all-nighter before a trial, never worried that a free AI chat tool might have just moved a client's confidential information somewhere outside the firm's control, and has never had to explain to The Florida Bar why an AI-generated research memo cited a case that does not exist.

Generic AI advice does not know three things that matter for a small law firm.

First, it does not know that your conversations are privileged. Every prompt that contains client information is governed by Rule 4-1.6, and the specific Claude plan you use changes your legal exposure under that rule. The distinction between a free consumer account and a paid enterprise account is not a product-marketing choice for a law firm. It is a confidentiality decision. Generic advice does not mention this because most of its audience does not care.

Second, it does not know about your bar rules. Florida Bar Ethics Opinion 24-1 is explicit: a lawyer using AI must take reasonable precautions to ensure compliance with the duties of confidentiality, supervision of nonlawyer assistance, candor, truthfulness, and fees. Generic AI advice was not written for professionals who face discipline, sanctions, and malpractice exposure for the outputs they rely on.

Third, it does not know your workflow. A small law firm is not a content-creation pipeline or a customer-service queue. Your work moves through identifiable stages, client intake, legal research, document drafting, review, client communication, administration, and business development, and each stage has a different AI risk profile, a different tool fit, and a different verification requirement. Generic advice treats all knowledge work as interchangeable. It is not.

The Cost of Wrong

What happens when a Florida attorney follows generic AI advice? She uses a personal Free account to draft a client document because the advice said "just start using it." Free, Pro, and Max users choose a model-improvement setting; a plan name alone does not tell you whether submitted conversations are used to improve Claude. If that setting is enabled, data in the improvement pipeline may be retained for up to five years. Even if it is disabled, the firm has put client information into an unapproved consumer account without checking its terms or controls. That creates serious Rule 4-1.6 risk before the draft is finished. [S1], [S2], [S7]

He relies on an AI research memo without verifying each citation against a primary source, because the advice said 'AI saves hours on research.' He files it. Dozens of cited authorities were fabricated or unsupported. The court strikes the brief, fines him, and refers him to the Bar. This happened in Florida. The case is Capital Standard, LLC v. U.S. Bank National Association, No. 2D2024-1392, 2026 WL 2450571 (Fla. 2d DCA Aug. 21, 2026).

She lets a paralegal use an AI tool to draft twenty engagement letters and save them to the client folder, no attorney review, because the advice said 'AI is just a starting point.' Six letters have the wrong matter type. One has the wrong client name. Two promise billing terms the firm never intended to offer. This is a Rule 4-5.3 supervision failure, and it scales fast when the tool can process twenty matters as easily as one.

What This Playbook Does Differently

This playbook was built for the Florida firm with fewer than 25 lawyers. Not for a marketing agency, not for a SaaS startup, not for a general small business. It covers two of the three Claude interfaces: Claude Chat and Claude Cowork. It answers three questions that matter:

1. What are the real ethics risks of using Claude for my work, and which guardrails will hold?
2. Which interface do I use for which tasks, and what is the difference between Chat and Cowork?
3. How do I implement Claude without creating a confidentiality breach, a sanctionable filing, or an ethics violation?

No evangelism. No predictions about the future of law. No vendor payments. Just the Florida-grounded analysis a small law firm needs to make good decisions about Claude in 2026.

The Actionable AI Philosophy

Tools serve lawyers. Not the other way around.

The firm that benefits most from Claude is not the one that adopts it fastest. It is the one that adopts it most deliberately, picks the right plan, governs the right tasks, and builds verification gates at the right points. Claude Chat handles one conversation at a time. Claude Cowork executes multi-step tasks autonomously. Both interfaces are available to you on a paid plan. The question is not whether to use them. The question is where each earns its keep and where it creates risk.

> **THE ACTIONABLE AI TEST**
>
> Before you put Claude near a client matter, answer four questions:
>
> 1. Which Claude plan am I on, and does it give me the data protections this matter requires?
> 2. Where does the client information go, and who can access it?
> 3. Can an attorney verify this output before it reaches a client, a court, or opposing counsel?
> 4. Would I be comfortable if The Florida Bar, and the judge or partner on this matter, knew I used Claude this way?
>
> If you cannot answer all four confidently, do not proceed.

Why Small Law Firms Need a Different AI Approach

Every business has its own relationship with AI. But before you get to any tool-specific guidance, there is a more fundamental challenge: small law firms have structural features that most AI advice ignores. Those features change which tools to use, how to govern them, and where the real risk lives.

The Non-Technical Partner Is the Decision-Maker

Most small law firms do not have an IT department. The managing partner who decides whether to adopt Claude is the same person who learned to type in law school, who still prefers printed briefs, and who has every right to be skeptical of AI for the right reasons. The two real objections, that AI makes things up and that AI might compromise client confidences, are not fear of change. They are accurate risk assessments that deserve specific answers, not reassurance.

Generic AI advice was written for technology workers who are comfortable reading documentation and setting up accounts. The implementation path for a small law firm has to work for someone who does not want to learn a new system. If the tool requires a steep learning curve, the partners will not use it, the paralegals will use it unsupervised, and the risk materializes without the benefit. The simplest tool that solves the problem is almost always the right answer.

Three Claude Interfaces, and What This Playbook Covers

Anthropic makes Claude available through three distinct interfaces. Understanding the difference is the first decision every small law firm needs to make.

Claude Chat is the conversational web and mobile interface available at claude.ai. You open a browser or the mobile app, type a question or a request, and Claude responds. It handles individual tasks well: a research question, a first-draft document, a client email, a quick analysis. Chat works in individual conversations, though a paid subscription adds Projects, persistent workspaces with custom instructions and uploaded files that give Claude context for your practice.

Claude Cowork is an agentic assistant in the Claude Desktop application for Windows and macOS and on supported web and mobile plans. Unlike Chat, Cowork does not wait for your next prompt. You give it a task, not a question, and it works through the steps: reading permitted files, creating documents, and organizing output. It became generally available on April 9, 2026. Its screen location and available features depend on plan, administrator settings, and surface. [S3]

> **THREE INTERFACES, TWO IN THIS PLAYBOOK**
>
> Claude Chat: The web and mobile conversational interface at claude.ai. Best for individual tasks, one conversation at a time.
>
> Claude Cowork: The agentic assistant in Claude Desktop on paid plans, with web/mobile access on supported plans (Enterprise requires admin enablement). Best for multi-step workflows. Local folder access requires desktop; computer use is Pro/Max only.
>
> Claude Code: The command-line interface tool for software developers. This playbook does not cover Claude Code.

A Note on Claude Code

Claude Code is Anthropic's command-line interface tool designed for software developers writing, reviewing, and debugging code. It is accessed through a terminal window and requires familiarity with developer tools. For the vast majority of small law firm attorneys, paralegals, and administrative staff, Claude Code is not the right tool, and this playbook does not cover it.

The key distinction: Claude Chat and Claude Cowork are built for knowledge work, reading, writing, analyzing, organizing, drafting, and communicating. Claude Code is built for programming. If you are not writing computer programs, Claude Code is not your interface.

One practical note for Pro subscribers: Claude's $20-per-month Pro plan includes access to Claude Code in the terminal. If you have a Pro subscription and notice a Code-related feature, it is safe to ignore it for day-to-day legal work. Your work belongs in Chat or Cowork.

Claude Chat: The Conversational Interface

Claude Chat is available on five plans in 2026. The plan you choose is the most important decision in your Claude adoption, because it determines what data protections you have.

— **Free ($0/month):** Basic access under Anthropic's consumer terms, including the model-improvement setting described below. Do not use the Free plan for any prompt that contains client information.

— **Pro ($20/month):** A consumer plan. Since Anthropic's September 2025 consumer-terms change, Free, Pro, and Max accounts have a model-improvement choice. If enabled, conversations in the improvement pipeline may be retained for up to five years. Anthropic's 2025 notice described 30-day retention for people who decline model improvement; its current privacy article separately says backend deletion occurs within 30 days after a user deletes a chat, subject to stated exceptions. Do not treat those statements as the same retention rule. Verify the current period, exceptions, and your account setting before use. Includes Projects, memory, and extended context. Suitable for work without client-identifiable information. [S1], [S2]

— **Max ($100–$200+/month):** A higher-usage consumer plan with expanded limits, in two tiers. Same consumer terms as Pro, including the model-improvement setting. Verify current features and pricing at anthropic.com.

— **Team (Standard seats $25/seat/month billed monthly, or $20/seat/month billed annually; higher-priced Premium seats are also offered; minimum 2 members, and the plan caps at 150 users):** A commercial plan. Anthropic's commercial terms do not use customer data to train models by default. Its standard Data Processing Addendum (DPA) is incorporated into the Commercial Terms; document the firm's acceptance and verify current data-handling terms before client work. Adds admin controls, centralized billing, SAML SSO, and workspace governance. The baseline for a firm using Claude on client matters. [S4]

— **Enterprise (published self-serve terms as checked September 29, 2026: $20/seat/month billed annually, plus metered usage at API rates; annual commitment and 20-seat minimum):** A commercial plan for advanced governance and security requirements. Sales-assisted arrangements may differ. The 20-seat minimum matters to firms with fewer than 25 lawyers; verify the full cost, feature coverage, and current terms. Enterprise alone does not establish HIPAA coverage for PHI. [S5], [S11]

Chat's Projects feature is worth specific attention for law firms. A Project lets you give Claude standing instructions, 'You are assisting a small Florida family law firm. Never disclose client names in an output. Always note when a citation needs verification.', along with uploaded reference files such as your firm's intake checklist, standard fee agreement, or legal research templates. Every conversation inside that Project starts with that context already loaded.

Claude Cowork: The Agentic Assistant

Claude Cowork is a fundamentally different relationship with AI than Chat. Chat answers your prompts one at a time. Cowork takes ownership of a task.

When you assign a task in Cowork, 'Review all documents in the client intake folder, identify any missing items from the standard checklist, and draft a missing-documents letter for each client', Cowork reads the files, applies the checklist, identifies the gaps, drafts the letters, and saves them to the folder. You review the output. The task ran while you were doing something else.

Cowork runs in Claude Desktop on every paid plan. Web and mobile access is on Pro, Max, and Team, and on Enterprise when an administrator enables it. The merged interface is rolling out, so controls may differ. Local folder access requires Desktop. As checked September 29, 2026, computer use is limited to Pro and Max; Team and Enterprise do not have it. Because this playbook recommends Team or Enterprise for client matters, do not build a client workflow on computer use. Multi-step tasks, schedules, and Skills depend on current features. [S3], [S10]

The governance implication of Cowork is material. An agent that can execute tasks autonomously needs explicit limits. Your governance policy must specify what Cowork can do without attorney review, what requires approval before execution, and what it must never do unsupervised.

The Confidentiality Imperative: Why Your Plan Choice Matters

Florida Bar Ethics Opinion 24-1 applies Rule 4-1.6's confidentiality duty to AI use. Consumer terms may permit provider access inconsistent with that duty. Check the plan's actual terms and controls before submitting client information.

The practical rule: no client identifiers in Free or Pro. For client matters, verify the Team or Enterprise Commercial Terms, the incorporated Data Processing Addendum, access controls, and the task's data needs. This playbook sets Enterprise plus attorney approval for Tier 4 as firm policy. For PHI subject to HIPAA, the Enterprise Primary Owner must activate HIPAA compliance and accept Anthropic's BAA; use only covered features. Cowork is excluded. The self-serve Enterprise offer has a 20-seat minimum. Account for that cost or keep the data out of Claude. [S4], [S5], [S11]

THE PLAN DECISION

Before you put any client information into Claude, Chat or Cowork, confirm which plan you are on and read the current terms of service.

Terms change. What was true six months ago may not be true today.

Current terms: anthropic.com/legal

For Team/Enterprise: verify the Commercial Terms and incorporated Data Processing Addendum, and record acceptance before client work. For PHI, confirm BAA activation and covered features.

The Florida Ethics Reality

The Florida Rules of Professional Conduct were not written with AI in mind. They apply anyway. The framework is not complicated once you see how the existing duties map onto Claude specifically.

Florida Bar Ethics Opinion 24-1, issued January 19, 2024, addresses generative AI directly and should be treated as essential reading for any Florida lawyer using AI tools on client matters. Read it in full at floridabar.org; the analysis below is not a substitute.

The Five Rules That Matter Most

RULE 4-1.1: COMPETENCE (AND TECHNOLOGICAL COMPETENCE)

The comment to Rule 4-1.1 requires lawyers to keep up with the benefits and risks of relevant technology. For Claude specifically, competence has two faces. You must understand what Claude can do, draft a first-cut document, synthesize a research question, organize a set of files, and what it cannot do reliably, which includes verifying that a cited case is real and accurately quoted, knowing the current version of a recently amended Florida statute, and exercising judgment on a matter that turns on facts it cannot see.

Competence also means understanding your plan's data protections. A lawyer who is using the Free tier for client matters without knowing what the terms of service say about data use is not competent under Rule 4-1.1.

RULE 4-1.6: CONFIDENTIALITY OF INFORMATION

Rule 4-1.6 requires that a lawyer shall not reveal information relating to representation of a client. The confidentiality obligation applies to everything you put into a Claude prompt. A prompt that includes a client's name, matter details, or any identifying information is disclosure of client information to a third party, the AI provider.

The firm policy recommended here is Team or Enterprise for client-specific work, with the current Commercial Terms and incorporated Data Processing Addendum verified and acceptance recorded. No client identifiers go into Free accounts. For PHI, the additional BAA and covered-feature conditions apply. This is a practical implementation of the confidentiality duty, not a claim that the standard DPA requires a separate signature. [S4], [S5], [S7]

> **THE CONFIDENTIALITY BRIGHT LINE**
>
> Client identifiers, Social Security numbers, financial account numbers, and sensitive personal information never go into a Free-tier Claude account.
>
> Client-specific work belongs on Team or Enterprise plans, with verification of current data-handling terms.
>
> For Tier 4, firm policy requires approved Enterprise use and attorney approval. For PHI subject to HIPAA, require an activated BAA and covered feature; Cowork is excluded. The published self-serve Enterprise offer has a 20-seat minimum.
>
> When in doubt, strip identifiers first, or do not paste it at all.

RULE 4-5.3: RESPONSIBILITIES REGARDING NONLAWYER ASSISTANCE

Florida Bar guidance applies the supervision principle to AI tools: you are responsible for Claude's output the same way you are responsible for a paralegal's work. In a small firm, the volume pressure creates the specific risk. It is tempting to let a paralegal run intake summaries, draft routine letters, or process documents through Cowork and pass the output forward without attorney review.

The supervision principles have special weight for Cowork. Opinion 24-1 applies standards for nonlawyer assistance by analogy to generative AI; the system itself is not a nonlawyer employee. When Cowork processes files, drafts letters, and saves output to a client folder, an unsupervised AI system is performing client-related work. The firm's policy must specify which outputs require attorney review before anything else happens with them. [S7]

RULE 4-3.3: CANDOR TOWARD THE TRIBUNAL

Claude can invent cases, misquote real opinions, and miss statutory amendments. Rule 4-3.3(a)(1) prohibits a lawyer from knowingly making a false statement of fact or law to a tribunal. An AI error in a filing is not automatically a candor violation or a competence failure. Separately, Rule 2.515(d)(2) requires a Florida state-court signer to represent that cited authorities exist and are accurately cited; false or unverified citations may lead to sanctions after notice and an opportunity to be heard. A lawyer's verification and response to discovered errors matter under the competence and candor duties. [S6], [S9]

Capital Standard, LLC v. U.S. Bank National Association, No. 2D2024-1392 (Fla. 2d DCA Aug. 21, 2026), is the Florida benchmark. Appellant's amended initial brief contained at least 31 fabricated or unsupported citations; after opposing counsel flagged the problem, a reply brief added 15 more. The court struck both briefs and imposed $1,500 in fines: $500 for a late show-cause response and $1,000 for AI misuse. Counsel was made personally liable for U.S. Bank's fees incurred preparing its answer brief in response to the amended initial brief, and the court referred him to The Florida Bar. Verify every authority against a primary source before filing. [S8]

RULE 4-1.5: REASONABLE FEES

AI tools compress time. When Claude drafts in twenty minutes what a paralegal would have taken two hours to draft, the billing question is real. Florida Bar Ethics Opinion 24-1 is explicit that AI may not be used to inflate billed time. Rule 4-1.5 requires fees to be reasonable. Bill the time actually spent, not the time a task used to take.

The corollary: charging separately for AI tool use without disclosure is problematic. If a client's matter generates AI-related costs, that should be addressed in the engagement letter. If it is the firm's overhead cost of doing business efficiently, it belongs there, not as a line item the client never agreed to pay.

Other Rules That Touch This Work

Rule 4-1.4 (Communication) requires keeping the client reasonably informed. Whether your firm uses AI on the client's matter is a question some clients will ask. Having a clear, honest answer, and ideally a disclosure statement in the engagement letter (see Appendix B), is both the ethical and the practical approach.

Rule 4-1.18 (Duties to Prospective Clients) applies when AI tools are used at intake before an engagement is established. Information provided by a prospective client in an intake questionnaire processed through an AI tool is still governed by the confidentiality duty.

The Courts Are Writing Their Own Rules: Right Now

Beyond the Rules of Professional Conduct, courts across Florida and the federal system are establishing their own AI requirements for filings. These rules change quickly. Check before every filing.

THE NEW STATEWIDE RULE

Florida Supreme Court Opinion SC2026-0673 (May 28, 2026, effective June 15, 2026) amended Rule of General Practice and Judicial Administration 2.515(d)(2).

Every attorney signing a document filed in a Florida state court must represent that cited authorities 'exist and are accurately cited.'

This applies to Florida state courts only. Federal courts, including all three of Florida's federal districts (Northern, Middle, and Southern), are governed by their own rules and individual judge standing orders.

Authorized sanctions: reprimand, contempt, stricken document, dismissal, costs, and attorney's fees.

The rule does not prohibit AI use. It requires that the signing attorney has verified the authorities.

Florida Supreme Court Administrative Order AOSC26-12 bars courts from imposing local AI-use disclosure or filing-accuracy certification requirements in Florida state courts through administrative orders, policies, or judicial practices. It does not bar AI education or warnings about sanctions. Rule 2.515(d)(2), as amended in SC2026-0673, governs the signer representation for Florida state-court filings; check any other applicable court instructions before filing. [S9], [S12]

In federal court, the Northern, Middle, and Southern Districts of Florida each have judges with individual standing orders governing AI use and citation verification. There is no single federal rule: check the assigned judge's standing orders before every federal filing.

BEFORE EVERY FILING

In a Florida state court: does SC2026-0673 apply? Follow Rule 2.515(d)(2) as amended, AOSC26-12 preempts circuit-level AI disclosure and certification requirements in Florida state courts.

In federal court in Florida: check the assigned judge's current standing orders. No assumptions.

In any court: verify every citation in the document against the actual source.

The Emerging Issue Most Small Firms Are Missing: From AI Chat to AI Agent

Most small law firms are treating AI as a search engine with a nicer interface. You type a question, read the answer, and move on. That is conversational AI, and it has real value. But it is also the 2024 version of AI for law firms. The shift that is happening now, the shift most small firms have not registered, is from AI that responds to AI that executes.

Claude Chat answers your prompts one at a time. You ask, it responds. Claude Cowork takes a task and works through it. You delegate, it executes. That is a different category of tool, with different implications for what a small firm can do and different governance requirements.

What an AI Agent Actually Does

Here is the practical difference. In Claude Chat, you type: 'Draft a follow-up email to a client about the status of their matter.' Claude generates a draft. You review it, personalize it, send it. One email, one task, done in a few minutes.

In Claude Cowork, you say: 'Review the open matters in the client intake folder, identify any with no activity note in the last thirty days, draft a brief status update email for each, and save the drafts to the client folder for my review.' Cowork reads the files, identifies the matters, drafts the emails, saves them to the right locations. You review the batch and send. A task that used to take an hour of paralegal time happens while you are in court.

This is not incremental improvement in the same tool. It is a different relationship between attorney and technology. The attorney delegates. The agent executes. The attorney reviews and takes responsibility for the output.

The Scale of the Shift

Claude Cowork became generally available on April 9, 2026 and ships on every paid plan, with capabilities still expanding. The firms that understand what agentic AI means and build governance around it now will have a structural advantage as those capabilities expand. The firms that are still thinking about AI as a 'fancy search engine' are already a generation behind.

The specific opportunity for a small law firm: a two-attorney firm running Cowork on consistent administrative workflows, intake processing, routine status communications, document organization, billing narrative drafts, content production, can sustain an operational capacity that used to require more staff. Not by replacing attorney judgment. By automating the work that never required attorney judgment but was consuming attorney time anyway.

Three Immediate Opportunities

— **Batch client communications:** Instead of drafting one status email at a time, Cowork can review the entire open-matter list, identify which clients are due for a touchpoint, draft the emails, and save them for review.

— **Intake document processing:** When a new client submits a packet, intake form, prior documents, financial records, Cowork can read the folder, extract key information against a checklist, flag missing items, and draft the intake summary memo.

— **Scheduled administration:** Cowork's scheduled task feature runs processes on a set cadence. A weekly deadline check across all open matters. A monthly reminder to review file status. A recurring task that drafts this week's billing narratives from the activity notes.

THE AGENT QUESTION

Before you assign a multi-step task to Cowork, ask: If Cowork completes this task exactly as I described, and I'm wrong about one assumption, what is the worst thing that happens?

If the answer is 'a draft I would review before anything else happens'. Run the task.

If the answer is 'an email that goes out without my review' or 'a file that reaches someone it should not'. Refine the task first.

The governance policy must specify: what tasks can Cowork run autonomously, and what requires attorney approval before the output moves forward.

The Governance Implication

Agentic AI creates a governance question that conversational AI does not: who is supervising the agent while the attorney is doing something else? The answer under Rule 4-5.3 is the same as it always was. The attorney is responsible for the output. But the practical mechanism of supervision changes when the output is a batch of fifty documents rather than a single draft on your screen.

Every firm that adopts Cowork needs an agent policy: a written set of rules that specifies which tasks Cowork can execute without real-time attorney oversight, which require pre-approval, which require post-review before anything is transmitted, and which tasks Cowork may not perform at all. The Sample AI Governance Policy in Appendix A includes an agent policy section as a starting point.

Ethics Risks and Governance Framework

Five specific failure modes are waiting for small law firms that adopt Claude without governance. Name them, because unnamed risks are the ones that land in your Bar complaint.

High-Risk Failure Modes

THE CONFIDENTIALITY BREACH VIA CONSUMER ACCOUNT

A paralegal uses a personal Free account to draft a client document because it is faster than logging into the firm account. The account's model-improvement choice has not been checked. If improvement is enabled, submitted conversations may enter a pipeline with retention of up to five years. Even if disabled, the firm has not approved the account or verified its controls. Client details have been submitted to a consumer service without the firm's review. That is a serious confidentiality risk under Rule 4-1.6; the firm must assess the facts and respond. [S1], [S2], [S7]

HALLUCINATED RESEARCH CITATIONS

Capital Standard, LLC v. U.S. Bank National Association, No. 2D2024-1392 (Fla. 2d DCA Aug. 21, 2026), is the case to use when someone says, "I checked the AI output." Counsel used generative AI as a primary research tool. The amended initial brief contained at least 31 fabricated or unsupported citations, and the reply brief added 15 after opposing counsel flagged the problem. Both briefs were struck. The $1,500 in fines consisted of $500 for a late show-cause response and $1,000 for AI misuse. Counsel was personally liable for U.S. Bank's fees to prepare its answer brief responding to the amended initial brief and was referred to The Florida Bar. Verification is the minimum. [S8]

UNSUPERVISED COWORK EXECUTION

A paralegal assigns a Cowork task that processes twenty intake packets, drafts twenty engagement letters, and saves them to the client folder. No attorney reads them before the letters go out the next morning. Six letters have the wrong matter type. One has a billing-rate error. Two include fee arrangements the firm never intended to offer. This is a Rule 4-5.3 supervision failure at production scale.

AGENT OVER-REACH

A Cowork scheduled task is set to send weekly status emails to all active clients. The 'send' step runs before the attorney reviews the drafts. One email references a settlement discussion the client asked not to be documented. Another was sent to a client whose matter was closed last week. Unsupervised transmission is the specific agent failure mode that generates Bar complaints and malpractice claims.

BILLING INFLATION

AI compresses a two-hour research task to twelve minutes. The time entry reads '2.0 hours, legal research.' Rule 4-1.5 requires fees to be reasonable. Florida Bar Ethics Opinion 24-1 specifically warns against using AI to inflate billed time. Bill the work actually performed at the time actually spent.

Data Classification Framework

Tier	What It Includes	Approved Claude Plans
Tier 1: Public	General legal research, public statutes, published	Any plan, including Free

Tier	What It Includes	Approved Claude Plans
	case law, publicly available information	
Tier 2: Anonymized	Matter facts with all client identifiers stripped, no name, no case number, no jurisdiction that narrows to the client	Pro or higher
Tier 3: Client Matter	Client identifiers, matter details, correspondence, and work product with client information	Team or Enterprise with current terms verified
Tier 4: Highly Sensitive	Social Security numbers, financial account numbers, medical or health records, immigration status, family conflict details	Enterprise, with current Commercial Terms and incorporated DPA verified; attorney approval per use. For PHI subject to HIPAA: Enterprise Primary Owner must activate HIPAA compliance and accept BAA; confirm covered feature. Cowork excluded. Published self-serve Enterprise minimum: 20 seats. [S4] [S5] [S11]

Approved Tool List and Governance Policy

Every firm using Claude needs a written AI governance policy. The policy should specify: which Claude plans are approved for which data tiers, who can authorize exceptions, what output types require attorney review before anything else happens, and what Claude may never do unsupervised. Appendix A provides a full sample policy as a starting point.

For multi-person firms: the Team plan's admin controls allow the firm administrator to set usage policies and monitor adoption. Use them.

Staff Training Requirements

Every person in the firm who uses Claude, attorney, paralegal, administrator, needs to know three things before they use the tool on any client matter:

— Which plan the firm uses, what data can and cannot go into it, and how to tell the difference

— That AI output is a first draft, every citation, every statute, every factual claim is verified before it goes anywhere

— That no AI output that will reach a client, a court, or opposing counsel leaves the firm without attorney review

Best Use Cases by Workflow Stage

A small law firm's work moves through identifiable stages regardless of practice area. Here is where Claude, Chat and Cowork, fits at each stage, and where it does not.

Stage 1: Client Intake and Intake Analysis

Intake is where AI delivers fast, visible value. The tasks are structured, the process is repetitive, and the output is internal, which means lower risk while you build confidence with the tools.

— **Claude Chat helps:** Summarize handwritten or dictated call notes into a structured one-page intake memo. Draft intake questionnaire language adapted to the matter type. Flag information gaps for follow-up before the first consultation.

— **Claude Cowork helps:** Process a folder of new intake documents, forms, prior records, correspondence, extract key information against the firm's intake checklist, flag missing items, and save a summary memo per client to the matter folder.

— **Neither does:** Tell you whether to take the client, catch the intuition about a difficult family dynamic, or make the conflict determination.

> **IMPORTANT**
>
> Intake information from a prospective client is covered by Rule 4-1.18 before any engagement is established.
>
> The same data classification rules apply at intake as they do during the matter.

Stage 2: Legal Research

Legal research is Claude's highest-value and highest-risk application for lawyers. The value is real: synthesizing a body of law, identifying the key cases and statutes on an issue, and producing a first-draft research memo in a fraction of the time. The risk is equally real: hallucinated citations, misquoted cases, and stale statutes presented with complete confidence.

— **Claude Chat helps:** Produce first-draft research memos on legal questions, with citations flagged for verification. Synthesize multiple cases into a plain-language summary of a legal standard. Identify the strongest and weakest arguments on both sides of an issue.

— **Claude Cowork helps:** Review a folder of uploaded research documents, synthesize the key points across them, and produce a consolidated memo. Extract and organize cited authorities from a stack of briefs or opinions.

— **Neither does:** Verify its own citations. Catch an amended statute that post-dates its training. Replace Westlaw, Lexis, or Fastcase for final citation verification.

> **CAREER-ENDING TRAP**
>
> The verification requirement is not a formality. Capital Standard, LLC v. U.S. Bank National Association, No. 2D2024-1392, 2026 WL 2450571 (Fla. 2d DCA Aug. 21, 2026), established the Florida benchmark: counsel who used generative AI as his primary research tool without verifying the

> citations drew fines, personal liability for U.S. Bank's fees to prepare its answer brief responding to the amended initial brief, and a referral to The Florida Bar.
>
> Build verification into the workflow: every citation confirmed in a primary database before the document is finalized.

Stage 3: Document Drafting

First-draft production is where AI delivers the most visible time savings. A document that took three hours to draft from a blank page takes thirty minutes to review and adapt from an AI first draft. The economics are real. So is the review requirement.

— **Claude Chat helps:** Draft demand letters, client letters, motion outlines, routine correspondence, and standard contract provisions from structured input. Adapt firm templates to a specific matter's facts. Produce multiple drafting alternatives for a contested clause.

— **Claude Cowork helps:** Batch-produce first drafts across multiple matters, multiple demand letters, multiple routine motions, multiple client updates, from a folder of matter data. Maintain consistent formatting and structure across a batch of related documents.

— **Neither does:** Know the specific judge's preferences, catch the procedural trap buried in the local rules, or make judgment calls about how hard to push.

Stage 4: Document Review and Quality Check

Review is where Cowork's batch capabilities are most powerful. Reviewing one document is fast either way. Reviewing a stack of documents is where AI changes the math.

— **Claude Chat helps:** Compare two contract versions and identify differences. Check a draft for internal consistency, do the defined terms match their use? Summarize a long document in plain English for a client.

— **Claude Cowork helps:** Review an entire folder of client-submitted documents against a completeness checklist. Flag documents that are missing, incomplete, or inconsistent. Produce a gap report across the batch.

— **Neither does:** Replace attorney judgment on whether a specific clause is acceptable. Catch the subtle ambiguity that will matter in litigation.

Stage 5: Client Communications

Client communication is where small firms lose hours every week: status updates, responses to routine questions, plain-English explanations of legal developments that affect the client. AI does not replace the relationship. It clears the production work so the attorney can focus on the relationship.

— **Claude Chat helps:** Draft status update emails from matter notes. Write plain-English explanations of legal developments, court orders, or settlement terms. Prepare meeting agendas and follow-up summaries.

— **Claude Cowork helps:** Review open matters across the firm's client folder, identify clients due for a status touchpoint, draft the status emails, and save them for attorney review.

— **Neither does:** Assess the emotional register of a difficult client relationship, navigate a conversation about bad news, or replace the attorney's judgment about what to say.

> **IMPORTANT**
>
> All client communications drafted by Claude, Chat or Cowork, are reviewed by the responsible attorney before they are sent.
>
> Cowork must not be set to send communications automatically. Draft, review, send.

Stage 6: Matter Administration

Administration is the category of work that consumes attorney and paralegal time without generating billed value: file organization, naming conventions, time-entry narratives, deadline tracking, and internal documentation.

— **Claude Chat helps:** Draft time-entry narratives from activity notes. Calculate deadline chains from a triggering date. Draft internal matter summaries for handoffs.

— **Claude Cowork helps:** Organize a client folder that has grown disorganized, rename files to a consistent convention, group them by type, produce a folder index. Run a weekly deadline check across all matter folders and produce an alert list.

— **Neither does:** Know which deadlines are hard statutory limits versus soft internal targets, or substitute for a docketing system with proper notification rules.

Stage 7: Marketing and Business Development

Marketing content is the lowest-risk AI use for a law firm. The output is not client-specific, the data is not confidential, and the worst outcome is a draft that misses the mark and needs revision.

— **Claude Chat helps:** Draft LinkedIn posts from a case insight, a bar update, or a client question you answered this week. Write blog post outlines and first sections. Produce email newsletter content.

— **Claude Cowork helps:** Maintain a content calendar, read the topic list in a folder, draft the month's posts, save them as individual files for review.

— **Neither does:** Know your clients' specific concerns better than you do, replace the insight that comes from years of practice, or make the final judgment call on what to publish under your name.

Prompts Library

Fifteen ready-to-adapt prompts, organized by workflow stage. Each is labeled with the interface, Chat or Cowork, it was written for. All are starting points: adapt the facts, add your firm's specific context, and verify every output before it moves forward.

BEFORE YOU USE THESE

Strip client identifiers, case numbers, and sensitive personal information before pasting into any tool. Match the data tier to the plan (see the Data Classification Framework).

Verify every citation, statute, and factual claim against a primary source before the output leaves the firm.

Every prompt produces a first draft. Every first draft is reviewed by an attorney before it goes anywhere.

Document AI use in the matter file.

Stage 1: Intake and Intake Analysis

PROMPT 01: NEW MATTER INTAKE SUMMARY MEMO · Stage 1: Intake | Claude Chat

```
You are a paralegal preparing an intake summary for a small law firm. From the
intake notes below, produce a one-page memo with the following sections:
1. Matter type and likely legal claims or issues
2. Key facts summary (3-5 sentences)
3. Parties: client, opposing party, any third parties
4. Documents received and documents still needed
5. Preliminary deadlines or urgency flags
6. Three open questions for the intake attorney
Do not give legal advice, predict outcomes, or recommend whether to take the
matter.
[PASTE ANONYMIZED INTAKE NOTES]
```

PROMPT 02: CONFLICT SCREEN ISSUES LIST · Stage 1: Intake | Claude Chat

```
You are assisting a small law firm with a preliminary conflict screen. From the
intake information below, produce a list of:
1. Names, entities, and relationships that should be run through the firm's
conflict database
2. Any facts that raise a potential conflict question under Rules 4-1.7
(concurrent conflicts) or 4-1.9 (former client conflicts)
3. Any relationships or prior representations mentioned that an attorney should
assess
Do not make the conflict determination. Present every item as a flag for attorney
```

```
review.
[PASTE ANONYMIZED INTAKE INFORMATION]
```

Stage 2: Legal Research

PROMPT 03: LEGAL RESEARCH SUMMARY MEMO · Stage 2: Research | Claude Chat

```
You are preparing a first-draft research memo for a small law firm on the legal question below. Produce a memo with:
1. Issue statement (one sentence)
2. Short answer (2-3 sentences)
3. Legal standard, cite the controlling rule, statute, or test with full citations
4. Case law analysis, summarize the key cases, with citations, in order of weight
5. Application to the facts presented
6. Open questions for further research
Flag every citation with [VERIFY]. Every case and statute must be confirmed in a primary legal database before this memo is used.
[PASTE THE RESEARCH QUESTION AND RELEVANT FACTS]
```

PROMPT 04: STATUTE INTERPRETATION AND APPLICATION · Stage 2: Research | Claude Chat

```
You are helping a small law firm understand a specific statutory provision and how it applies to a set of facts. From the statute and facts below, produce:
1. Plain-English explanation of what the statute requires or permits
2. How the provision has been interpreted by Florida courts, cite any relevant cases
3. How the statute applies to the specific facts presented
4. Any ambiguities, exceptions, or unsettled questions in the application
Flag every statutory citation and case citation with [VERIFY]. Note any effective date or recent amendment.
[PASTE THE STATUTORY PROVISION AND ANONYMIZED FACTS]
```

PROMPT 05: OPPOSING ARGUMENT ANTICIPATION · Stage 2: Research | Claude Chat

```
You are helping a small law firm prepare for the strongest arguments opposing counsel is likely to make on the issue below. Produce:
1. The three to five strongest arguments the opposing party is likely to raise
2. The authorities most likely to be cited in support of each argument
3. The best response to each argument, with supporting authority
4. Any argument that the firm currently lacks a strong answer to
Flag every authority with [VERIFY]. Present the analysis for attorney decision-making.
[PASTE THE LEGAL ISSUE AND THE FIRM'S CURRENT POSITION]
```

Stage 3: Document Drafting

PROMPT 06: DEMAND LETTER FIRST DRAFT · Stage 3: Drafting | Claude Chat

```
You are drafting a demand letter for a small law firm. From the matter details
below, produce a professional demand letter with:
1. Statement of representation
2. Summary of the facts supporting the client's position
3. Legal basis for the demand, with relevant authority [flag all citations VERIFY]
4. The specific demand, what is being requested and by what date
5. Consequences of non-response or rejection
6. Invitation to contact the firm to discuss resolution
Do not include client identifiers in the output. Use placeholders in brackets.
[PASTE ANONYMIZED MATTER DETAILS]
```

PROMPT 07: MOTION OUTLINE AND STRUCTURE · Stage 3: Drafting | Claude Chat

```
You are helping a small law firm structure a motion for the court proceeding
described below. Produce:
1. Recommended motion type and procedural basis
2. Proposed section structure with headings
3. One-sentence summary of the argument in each section
4. Key cases and statutes to anchor each argument, flag all with [VERIFY]
5. Any procedural requirements or page limits to confirm before filing
Do not write the full motion. Produce the structure and argument map for attorney
development.
[PASTE ANONYMIZED MATTER DETAILS AND THE RELIEF SOUGHT]
```

PROMPT 08: CONTRACT CLAUSE REVIEW AND ALTERNATIVES · Stage 3: Drafting | Claude Chat

```
You are helping a small law firm review and improve a contract clause. For the
clause below, produce:
1. Plain-English explanation of what the clause currently does
2. Risks or ambiguities in the current language
3. Two or three alternative versions: one more favorable to the drafter, one more
neutral, one more favorable to the counterparty
4. Any standard industry language or market practice relevant to this type of
clause
Present the alternatives for attorney decision. Flag any legal standard that
requires verification.
[PASTE THE CONTRACT CLAUSE AND IDENTIFY WHICH PARTY THE FIRM REPRESENTS]
```

Stage 4: Document Review

PROMPT 09: DOCUMENT COMPARISON AND GAP ANALYSIS · Stage 4: Review | Claude Chat

```
You are helping a small law firm compare two versions of a document. From the documents below, produce:
1. A section-by-section comparison of the two versions, what changed, what was added, what was removed
2. Any provision in the revised version that is materially different in effect, not just in wording
3. Any provision removed from the original that the attorney should decide whether to restore
4. Three to five questions the attorney should address before signing or submitting the revised version
Do not recommend whether to accept the changes.
[PASTE DOCUMENT VERSION 1 AND DOCUMENT VERSION 2]
```

PROMPT 10: INTAKE DOCUMENT COMPLETENESS REVIEW · Stage 4: Review | Claude Cowork

```
COWORK TASK: Review the files in the [CLIENT INTAKE FOLDER] against the firm's standard intake checklist below. For each client folder in the directory:
1. Identify which checklist items are present (with filename)
2. Identify which checklist items are missing
3. Flag any document that appears incomplete or unsigned
Produce a summary report organized by client folder, saved to the intake folder as 'Intake-Review-[DATE].docx'. Do not move or modify any client files. Read only.
[PASTE THE FIRM'S STANDARD INTAKE CHECKLIST]
```

Stage 5: Client Communications

PROMPT 11: MATTER STATUS UPDATE EMAIL DRAFT · Stage 5: Client Comms | Claude Chat

```
You are drafting a client status update email for a small law firm. From the matter notes below, produce a brief, professional email with:
1. A two-sentence summary of where the matter stands
2. What happened since the last update (if noted)
3. What happens next, the next step and the expected timing
4. Anything the client needs to do or decide
5. An invitation to call or email with questions
Keep it under 200 words. Plain English, no legal jargon. Do not predict outcomes.
[PASTE ANONYMIZED MATTER NOTES]
```

PROMPT 12: PLAIN-ENGLISH CASE EXPLANATION · Stage 5: Client Comms | Claude Chat

```
You are helping a small law firm explain a legal concept, court ruling, or procedural development to a client who is not a lawyer. From the information below, produce:
1. A plain-English explanation of what happened or what the term means, no jargon
2. What it means for the client specifically
3. What the next steps are
4. One or two questions the client is likely to ask, and the honest answers
Write at a level a non-lawyer can read in two minutes. Do not give legal advice or predict outcomes.
[PASTE THE LEGAL CONCEPT OR RULING AND HOW IT RELATES TO THE CLIENT'S SITUATION]
```

Stage 6: Matter Administration

PROMPT 13: TIME ENTRY NARRATIVE DRAFTS · Stage 6: Administration | Claude Chat

```
You are drafting time-entry narratives for a small law firm from activity notes. For each activity below, produce a professional billing narrative that:
1. Describes the work performed accurately and specifically
2. Uses active voice and clear language
3. Reflects only the time actually spent (do not inflate or estimate)
4. Does not disclose privileged strategy or confidential communications in the narrative
Produce one narrative per activity. Do not include client identifiers. Use matter codes or placeholders.
[PASTE ACTIVITY NOTES WITH TIME SPENT]
```

Stage 7: Marketing and Business Development

PROMPT 14: LINKEDIN POST FROM PRACTICE INSIGHT · Stage 7: Biz Dev | Claude Chat

```
You are drafting a LinkedIn post for a small law firm attorney. From the insight, question, or experience below, produce:
1. Three hook options (first line only, the one that earns the scroll)
2. One full post draft (250-350 words) that leads with the problem or insight, explains why it matters for small law firms or their clients, and ends with a clear takeaway or invitation to connect
Voice: direct, peer-to-peer, no jargon. The insight should be general enough to publish publicly.
[PASTE THE INSIGHT, CASE PATTERN, OR LEGAL DEVELOPMENT TO BUILD THE POST AROUND]
```

PROMPT 15: MONTHLY NEWSLETTER SECTION DRAFT · Stage 7: Biz Dev | Claude Cowork

```
COWORK TASK: Read the topic notes and research files in the [NEWSLETTER FOLDER] for this month. For each topic file present:
1. Draft a newsletter section of 150-200 words, plain English, one key takeaway, practical for a small business or individual client to read
2. Save each section as a separate file named 'Newsletter-Section-[TOPIC]-[DATE].md' in the newsletter folder
3. Produce an index file listing all sections drafted, with a one-line summary of each
Do not include client names or confidential information. Treat all sections as drafts for attorney review before publication.
[CONFIRM THE NEWSLETTER FOLDER PATH AND TOPIC NOTES ARE IN PLACE BEFORE RUNNING]
```

Appendix A: Sample AI Governance Policy

[Firm Name]: Artificial Intelligence Use Policy

Version 1.0 | Effective Date: [Date] | Review Date: [Date + 1 Year]

1. Purpose

This policy governs the firm's use of Claude Chat and Claude Cowork in connection with client matters, business operations, and marketing. It exists to protect client confidentiality, ensure compliance with the Florida Rules of Professional Conduct and Florida Bar Ethics Opinion 24-1, and establish clear accountability for AI-assisted work product.

2. General Principles

— AI is a tool. Lawyers supervise tools. The firm applies the supervision and review principles in Opinion 24-1 to every Claude output used in a client matter.

— No AI output goes to a client, a court, or opposing counsel without attorney review.

— The client's confidentiality interest governs every tool-selection and data-handling decision.

— If you are unsure whether a task or a prompt complies with this policy, stop and ask [Policy Supervisor name/title].

3. Definitions

— **Claude Chat:** The conversational AI interface available at claude.ai (web and mobile).

— **Claude Cowork:** The agentic assistant in Claude Desktop for paid plans; web/mobile access on Pro, Max, Team, and Enterprise where an admin enables it. Local folder access requires desktop, and computer use currently is Pro/Max only. [S3], [S10]

— **Approved Plan:** The firm's Claude plan for client-matter work: Team or Enterprise with current Commercial Terms and incorporated Data Processing Addendum verified and acceptance recorded. A separate signed instrument is required only if firm policy or negotiated terms demand it. Pro and Free are not approved for client-identifiable information. PHI needs the BAA activation and eligible-feature conditions below. [S4], [S5]

— **Agent Policy Supervisor:** The attorney designated to approve Cowork task configurations and review agentic workflow outputs.

4. Permitted Uses

— Legal research first drafts (with mandatory citation verification before any use)

— Document drafting assistance (with attorney review before finalization)

— Client communication drafts (with attorney review before transmission)

— Intake processing and document review (with attorney review of output)

— Administrative tasks, file organization, billing narrative drafts, scheduling assistance

— Marketing content, posts, newsletter sections, presentation materials (with attorney review before publication)

5. Prohibited Uses

— Any prompt containing client identifiers in a non-approved plan or non-approved account

— Tier 4 data in any account outside the firm-approved Enterprise configuration; for PHI subject to HIPAA, any use before the Primary Owner activates HIPAA compliance, accepts Anthropic's BAA, and confirms the feature is covered. Cowork may not process PHI under current BAA coverage. Attorney approval per use is required for Tier 4. [55]

— Cowork tasks configured to send communications without attorney review

— Cowork tasks that access files outside the specifically designated firm folder

— Using AI-generated research without independent citation verification in a court filing

— Using AI to generate billing entries that do not accurately reflect time actually spent

6. Confidentiality and Data Handling

No client identifiers may be included in any Claude prompt on a Free-tier account. Client-specific work is performed only on the firm's Approved Plan. Before using the Approved Plan for Tier 3 or Tier 4 data, verify the current terms of service at anthropic.com/legal.

7. Claude Cowork: Agent Policy

Cowork may execute the following tasks without real-time attorney oversight, subject to attorney review of output before any external use:

— Intake document processing and completeness review

— File organization and folder indexing tasks

— Billing narrative drafts from activity notes

— Marketing content drafts saved to the marketing folder

Cowork requires attorney pre-approval and post-review before any output is transmitted or used in a client matter:

— Any task involving client-specific documents

— Any task that produces draft communications intended for clients or opposing parties

— Any task involving Tier 3 or Tier 4 data

Cowork is prohibited from:

— Sending any external communication automatically, including email

— Accessing files outside the specifically designated firm folder

— Making scheduling decisions that could create a conflict

— Executing any action that cannot be reviewed and reversed before it affects a client or the firm

8. Court and Tribunal Filings

Any document filed in a Florida state court on or after June 15, 2026 must comply with Florida Supreme Court Opinion SC2026-0673. The signing attorney represents that every cited authority exists and is accurately cited. Citation verification against a primary source is mandatory before signing any filing that includes AI-assisted research.

Federal court filings are not governed by SC2026-0673. They are governed by the assigned judge's current standing orders and applicable local rules. Check your specific judge's standing orders before every federal filing.

9. Incident Response

If a team member suspects that client confidential information was disclosed through a Claude account that was not approved for that data, the team member immediately notifies [Policy Supervisor name] and stops using the account. The Policy Supervisor assesses the disclosure, determines whether a notification obligation exists, and documents the response.

10. Policy Acknowledgment

Each attorney and staff member who uses Claude confirms in writing that they have read and understand this policy. Annual recertification is required.

Appendix B: Client AI Disclosure Statement

The following language is suitable for inclusion in the firm's engagement letter. Adapt to your engagement letter style and have a Florida-licensed attorney review before use.

SAMPLE DISCLOSURE LANGUAGE

Use of Artificial Intelligence: Our firm uses AI-assisted tools, including Anthropic's Claude platform, to assist with certain aspects of legal work, including document drafting, research assistance, internal matter summaries, and client communications. AI-generated work product is always reviewed by a licensed attorney before it is used in your matter or transmitted to any party.

We use AI tools only on plans with data-handling protections appropriate to client information. We do not use AI to process your Social Security number, financial account numbers, or other highly sensitive personal data without your specific authorization and additional safeguards.

If you have questions about our use of AI in your matter, or if you prefer that AI tools not be used, please let us know.

Implementation notes:

— Review and update the disclosure annually as your tool stack changes.

— If a client opts out of AI use, document the opt-out in the matter file and notify team members working on the matter.

— Retain signed engagement letters with the AI disclosure as part of the matter file.

Appendix C: Court AI Disclosure and Certification

Use the appropriate certification based on the court where the document is filed. Verify the current applicable order before every filing.

SAMPLE AI REPRESENTATION LANGUAGE: FLORIDA STATE COURTS (SC2026-0673)

Consistent with Florida Rule of General Practice and Judicial Administration 2.515(d)(2), the undersigned certifies that [he/she/they] has reviewed all authorities cited in this [document type] and that each cited authority exists and is accurately cited.

Note: SC2026-0673 creates a representation built into the signing attorney's signature, not a separately mandated certification block. Including this language explicitly is a best practice. This applies to Florida state courts only; it takes effect for documents signed on or after June 15, 2026. Federal courts are governed by their own rules and the assigned judge's standing orders.

SAMPLE DISCLOSURE: AI-ASSISTED DRAFTING (WHERE REQUIRED OR CHOSEN)

The undersigned discloses that artificial intelligence tools were used to assist in the preparation of this [document type]. The undersigned has reviewed all AI-generated content, verified all citations against primary sources, and takes full responsibility for the accuracy of the content as filed.

Note: AOSC26-12 bars local AI-use disclosure requirements in Florida state courts. Federal courts and courts in other jurisdictions may impose their own requirements. Check the applicable rules and orders before deciding whether to include this language.

Implementation notes:

— Assign one attorney in the firm to maintain a current log of AI-related orders for each court where the firm regularly practices.

— Confirm whether disclosure is required or appropriate in the court before using the sample. AOSC26-12 bars local AI-use disclosure requirements in Florida state courts; federal courts follow their own rules and orders.

— For federal courts: check the Northern District of Florida (flnd.uscourts.gov), the Middle District of Florida (flmd.uscourts.gov), and the Southern District of Florida (flsd.uscourts.gov) websites for judge-specific standing orders.

Appendix D: Letter to AI Vendors: Nine Questions

Use this questionnaire when evaluating any AI vendor for firm use, or when upgrading from a standard consumer plan to a business or enterprise plan.

1. Do your terms of service or data processing agreement prohibit you from using client data submitted through our account to train or improve your AI models? Please provide the relevant contractual language.
2. Who has access to the content of our prompts and conversations, technical staff, reviewers, contractors? Under what circumstances and for what purposes?
3. How is our data stored, and where? Is it stored in the United States? Do you use subprocessors who have access to client data, and if so, who are they?
4. What encryption is applied to our data in transit and at rest?
5. What is your process for notifying us of a data breach that involves our account data, and what is your target notification timeframe?
6. Is your standard Data Processing Addendum incorporated into our Commercial Terms? Provide the current text and effective terms. If our firm requires a separately executed or negotiated instrument, will you sign one?
7. How do you handle a request by us to delete our account data? How quickly is deletion completed, and do you provide confirmation?
8. Has your platform undergone a third-party security audit or penetration test in the last twelve months? Can you provide the audit report or a summary?
9. What is your process for notifying us of material changes to your data-handling policies or terms of service, and how much advance notice do you provide?

HOW TO READ THE ANSWERS

A vendor that refuses to provide written answers to these questions is not a vendor for client matters.

Vague answers about ‘industry-standard’ practices without specifics are not answers.

The contractual no-training clause (Question 1) is the minimum threshold for using the tool with client-specific information.

File the answers, accepted Commercial Terms and incorporated DPA, and any separately executed instrument in the firm’s AI governance records.

Appendix E: 90-Day Implementation Roadmap

This roadmap assumes your firm is starting from a baseline of informal, individual AI use and building toward a governed, firm-wide implementation.

Days 1-30: Foundation

— Designate an AI Policy Supervisor, the attorney who will own the governance policy, answer team questions, and review Cowork task configurations.

— Audit current AI use: who on the team is using Claude now, on which plan, for what tasks? Identify any current use that does not meet the data requirements in this playbook.

— Confirm or establish the firm's Claude plan. Every person doing client work should be on the firm's approved plan, not a personal account.

— Draft and distribute the AI Governance Policy (see Appendix A) and hold a 30-minute team training on the three things every team member needs to know.

— Add the Client AI Disclosure Statement (Appendix B) to the firm's engagement letter template.

— Pick one low-risk workflow, intake summaries, billing narratives, or status email drafts, and run it through Chat for 30 days.

Days 31-60: Proficiency

— Review the first 30 days. What worked? Where did team members want to use AI but hesitated because of the policy? Adjust.

— Introduce Claude Chat's Projects feature: set up one Project for the firm's primary practice area with standing instructions and uploaded reference materials.

— Run your first Cowork task: a low-stakes batch task such as organizing a folder, producing a batch of billing narratives, or reviewing an intake folder for completeness.

— Add the appropriate court certification language (Appendix C) to the firm's filing checklist.

— Verify the current terms of service for the firm's Claude plan against the data handling described in this playbook.

Days 61-90: Optimize

— Expand the workflows using AI to two or three additional stages from the Best Use Cases section.

— Set up one Cowork scheduled task: a weekly intake folder review, a recurring billing narrative batch, or a monthly content-production task.

— Measure: how many hours per week is the team saving? Which tasks are returning the most value?

— Schedule an annual AI policy review date and add it to the firm calendar now.

— Send the Vendor Questionnaire (Appendix D) to Anthropic for the firm's current plan and file the response in the governance records.

> **WHAT SUCCESS LOOKS LIKE AT DAY 90**
>
> Every team member is using Claude on the firm's approved plan, not a personal account.

Every citation in every AI-assisted filing has been verified against a primary source.

The Policy Supervisor has reviewed and approved every Cowork task configuration.

The client engagement letter includes the AI disclosure.

The team has identified two or three workflows where AI saves meaningful time every week.

Sources and Ethics Anchors

— S1. Anthropic, Updates to Consumer Terms and Privacy Policy (Aug. 28, 2025): https://www.anthropic.com/news/updates-to-our-consumer-terms

— S2. Anthropic Privacy Center, How long do you store my data? (July 1, 2026): https://privacy.claude.com/en/articles/10023548-how-long-do-you-store-my-data

— S3. Claude Help Center, Get started with Claude Cowork, Availability: https://support.claude.com/en/articles/13345190-get-started-with-claude-cowork

— S4. Anthropic Privacy Center, How do I view and sign your Data Processing Addendum?: https://privacy.claude.com/en/articles/7996862-how-do-i-view-and-sign-your-data-processing-addendum-dpa

— S5. Claude Help Center, Business Associate Agreements for Commercial Customers: https://support.claude.com/en/articles/8114513-business-associate-agreements-baa-for-commercial-customers

— S6. The Florida Bar, Rules Regulating The Florida Bar, Chapter 4 (June 2026): https://www-media.floridabar.org/uploads/2026/06/2026_12-JUNE-Chapter-4-RRTFB-1.pdf

— S7. The Florida Bar, Ethics Opinion 24-1 (Jan. 19, 2024): https://www.floridabar.org/etopinions/opinion-24-1/

— S8. Second District Court of Appeal, Capital Standard v. U.S. Bank, No. 2D2024-1392 (Aug. 21, 2026): https://flcourts-media.flcourts.gov/content/download/2494161/opinion/Opinion_2024-1392.pdf

— S9. Supreme Court of Florida, In re Amendments to Rule 2.515, No. SC2026-0673 (May 28, 2026): https://flcourts-media.flcourts.gov/content/download/2489374/opinion/Opinion_SC2026-0673.pdf

— S10. Claude Help Center, Let Claude use your computer in Cowork, Current limitations: https://support.claude.com/en/articles/14128542-let-claude-use-your-computer-in-cowork

— S11. Anthropic, Claude Enterprise Plan, Get the Enterprise plan: https://claude.com/solutions/enterprise

— S12. Supreme Court of Florida, Administrative Order AOSC26-12 (May 28, 2026): https://flcourts-media.flcourts.gov/content/download/2489379/29182858

— In re Amendments to Rules Regulating The Florida Bar – Chapter 4, 393 So. 3d 137 (Fla. 2024) (SC2024-0032, eff. Oct. 28, 2024): amending AI-related comments to the professional rules.

About Actionable AI

Greg Yates is the founder of Actionable AI, a solo AI consulting practice based in Daytona Beach, Florida. Before launching Actionable AI, Greg spent his career as a corporate attorney, including as a partner at a major New York law firm, followed by work in career counseling and law school career services. He comes to AI consulting as a former practitioner, not a technologist. His clients are small law firms, and his measure of success is simple: a managing partner who walks away saying, 'This is not as hard as I thought it was, and it is going to make a real difference.'

Actionable AI is a solo AI-adoption consultancy for small law firms. Founded by Greg Yates, a former Big Law corporate equity partner at Steptoe (NYC Managing Partner) and Seyfarth Shaw, it helps small firms evaluate, select, and implement AI tools without the learning curve, the sales deck, or the vendor agenda. Every recommendation is grounded in the specific economics, ethics obligations, and workflow of a non-technical small firm.

What We Do

— AI readiness assessments with specific, written findings and numbered ROI estimates you can hold us to

— Tool selection and implementation, picking the right tools and the right plan for your firm's data and workflow

— Staff training, hands-on, practical, no jargon

— Governance policy drafting, a written policy your team will actually follow

Who This Playbook Is For

This playbook is for the Florida firm with fewer than 25 lawyers that uses, or is considering using, Claude Chat or Claude Cowork. It was not written for a tech company, a large firm with an IT department, or an AI enthusiast. It was written for the managing partner who wants to make good decisions about these tools without becoming a technology expert.

READY TO IMPLEMENT CLAUDE IN YOUR FIRM?

Book a free 15-minute assessment call to discuss (https://actionableai.s.gy/9Ei2uw):

Which Claude plan is right for your firm's data requirements and headcount

Which two or three workflows to start with given your practice area and case mix

How to build a governance policy your team will actually follow

The Florida ethics rules and filing requirements that govern your work

Whether your current AI use, formal or informal, has gaps that need to be closed

No pitch. No deck. No obligation.

actionableai.net | Greg Yates | gyates@actionableai.net

Actionable AI

Download the free Florida AI Starter Playbook.

Eight practical resources for managing partners and firm administrators: a readiness checklist, five pilot workflows, a tool-approval scorecard, a lawyer verification checklist, and a 90-day plan. It is free, and the button below downloads it directly.

Scan to download the Starter Playbook

Download the Starter Playbook (PDF)

Direct download: Florida AI Starter Playbook (PDF)

Want a second set of eyes on which workflow to pick?

I take 15-minute calls with managing partners and administrators. I will tell you whether there is a fit. No deck.

Book a 15-minute introductory call

Book a 15-minute introductory call