



VdVLaw
Solutions Beyond Boundaries

AI Safety Toolkit

For Attorneys

*Internal use rules, client intake questions, client scripts,
and a handout clients can actually understand*

Prepared by Robert van der Vijver, Case Cartographer, VdVLaw
VdVLaw | www.vdvlaw.net
Current as of October 2, 2026

Attorney-directed legal support. Not a law firm. No legal advice.

About this toolkit

This toolkit gives attorneys and law office staff a practical, plain-language system for using generative AI without creating privilege, confidentiality, protective-order, preservation, filing, or verification problems. It reflects legal developments, ethics guidance, and vendor policies reviewed through October 2, 2026. This area changes quickly, so settings, contracts, rules, and court orders should be checked again before relying on them.

What is inside

- Part 1: Attorney AI safety rules, from account types to California's new statutory requirements.
- Part 2: Client intake questions, including third-party AI use after the client discussed the matter with someone else.
- Part 3: Client scripts, including a full version, a short engagement-letter version, and firm AI-use disclosure language.
- Part 4: A small client handout written for actual clients.
- Part 5: Sample office policy language and a California filing checklist.
- Part 6: Current case and rule examples attorneys should know.
- Part 7: Sources reviewed.
- Companion form: Client AI Use Intake Form and Acknowledgment with Attorney AI Consent.

October 2026 legal update

California SB 574 was signed on September 30, 2026 and takes effect January 1, 2027. It adds Business and Professions Code section 6068.1 and amends Code of Civil Procedure section 128.7. The toolkit now separates this enacted statute from the State Bar's still-proposed AI-related Rules of Professional Conduct amendments. It also adds the Ninth Circuit's *Lnu v. Blanche* discipline order, California's published *Del Biaggio v. Bansen* decision, and the updated status of the *New York Times v. OpenAI* preservation dispute.

Plain-English disclaimer

This document is a practical training tool for attorneys and law offices. It is not legal advice, does not replace jurisdiction-specific ethics review, and does not decide privilege, work product, protective-order, confidentiality, or statutory-compliance questions for any specific case. Lawyers should check applicable rules, statutes, orders, client obligations, vendor terms, and actual workspace settings before using AI with case information.

The short version attorneys should remember

AI is useful. It is not magic. It is not your lawyer. It is not your client. It is not your diary. And it is definitely not a privilege force field with a blinking cursor.

Use it carefully and it can help with organization, drafting, issue spotting, summaries, checklists, and workflow. Use it casually with client facts, discovery, medical records, police reports, juvenile records, protective-order material, or litigation strategy, and you may have created a new problem with better grammar.

Core rule

Not used for model training is not the same thing as privileged, confidential, protected by a court order, statutorily compliant, or safe to upload.

Part 1: Attorney AI safety rules

1. Know which account you are using before you type anything

There is a meaningful difference between a consumer account, a business or enterprise workspace, an API setup, and a legal-specific AI platform. The label on the login screen matters. So do the terms, retention settings, admin controls, connectors, sharing options, and whether the tool uses inputs or outputs for model improvement.

For public marketing, first drafts, general outlines, and nonconfidential research, AI can be a useful assistant. For client facts, litigation strategy, discovery, medical records, criminal records, juvenile records, sealed material, or protective-order material, the responsible attorney should approve the specific tool, workspace, settings, and matter use before case data is entered.

Firm standard: client work happens only in the firm's configured commercial workspace: ChatGPT Business or Enterprise, Claude for Work Team or Enterprise, Google Workspace with Gemini, or Microsoft 365 Copilot on the firm tenant, each configured and verified for the intended use. Consumer, free, and personal accounts are prohibited for client work. DeepSeek's hosted services are prohibited for firm work at any tier. If the correct commercial subscription and controls are not in place, the answer is not better prompting. It is stop.

Do not confuse plan name with compliance

A commercial plan can provide important contractual and technical safeguards, but SB 574 does not create a brand-name safe harbor. For California attorneys, the statutory text focuses on whether access to confidential, personal identifying, and other nonpublic information is restricted to the attorney and persons the attorney authorizes under confidentiality obligations. The actual workspace, access, integrations, sharing, and contract still matter.

2. Separate the questions: training, privilege, confidentiality, access, retention, and court orders are not the same thing

A tool may say it does not train on business data by default. Good. That answers one question. It does not answer the rest.

- Training: Will the provider use prompts, files, outputs, or feedback to improve models?
- Access: Who can access prompts, files, outputs, logs, workspace history, connected sources, and administrative exports?
- Retention: How long is the material retained in product history or backend systems, and what settings or legal holds change that?
- Privilege and work product: Does the use fit within applicable protections, and has counsel used reasonable safeguards?
- Protective orders and sealing: Was the lawyer or client allowed to place the material into the system at all?
- Discovery: Could prompts, outputs, uploads, logs, or the identity of the platform become discoverable or subpoenaed?

- Statutory compliance: For California matters, does the use satisfy Business and Professions Code section 6068.1 when operative, including restricted access, verification, correction, and court disclosure?

3. Treat AI prompts like case notes that may someday have to be explained

If a prompt contains client facts, case theory, witness analysis, or discovery content, do not treat it like vapor. Treat it like a record. The delete button is not a legal doctrine.

- Do not put raw confidential facts into a prompt unless the tool, workspace, settings, and use have been approved.
- Use placeholders for work that does not require identification: Client A, Officer 1, Witness 2, Incident Date 1, County A.
- Remove names, dates of birth, addresses, phone numbers, emails, medical identifiers, juvenile details, case numbers, and other unnecessary identifiers when the task does not require them.
- Avoid prompts that unnecessarily reveal strategy unless the environment and matter use have been approved.
- When AI helps generate work product, verify the output against the source record. The machine does not get to be the final witness.
- Record enough about approved use to reconstruct what system was used, what category of material went in, and what verification occurred, without creating unnecessary narrative about strategy.

4. Use the red, yellow, green test

Green zone	Yellow zone	Red zone
Generally safer: public-source research, nonconfidential drafting, marketing copy, general checklists, public statutes and cases, and internal training prompts with no client facts.	Requires attorney review: anonymized hypotheticals, redacted summaries, placeholder drafts, approved commercial workspaces, and work governed by a written office policy.	Do not upload outside the firm's configured commercial workspace, and never without matter-specific review: raw discovery, police reports, bodycam, jail calls, phone extractions, medical or juvenile records, sealed or protective-order material, client strategy, admissions, private photos, or confidential settlement communications.

Red-zone update: raw discovery is not automatically forbidden if the firm has a configured commercial workspace and the responsible attorney has completed matter-specific review. The attorney should check client authorization, protective orders, sealing restrictions, regulated data, access controls, connectors, retention, and deletion or preservation duties before upload. For California attorneys, SB 574 adds an additional restricted-access requirement when section 6068.1 becomes operative.

5. Do not let staff create shadow AI workflows

The danger is not always the lawyer. Sometimes it is the well-meaning assistant, investigator, intern, contract paralegal, expert, vendor, or client using a random AI site because it gave them a neat table once and now they trust it more than gravity.

- Maintain a written list of approved tools and approved account types.
- Ban unapproved browser extensions, third-party GPTs, agents, plug-ins, connectors, and file-upload tools for case materials.
- Require staff, contractors, experts, and investigators to ask before using AI on case-specific material.
- Train everyone that “I only pasted a little bit” is not a security policy.
- Use MFA, strong passwords, role-based workspace permissions, and separate personal and professional accounts.
- Ask who actually performed the AI step. A written policy does not answer that question after the fact.

6. Verify everything before it leaves the office

AI can draft with the confidence of a partner and the accountability of a vending machine. The lawyer still has to check the work.

- Check every case citation, quotation, statute, rule, and factual proposition against a reliable source.
- Do not rely on an AI summary of a source document without verifying important statements against the source page, timestamp, exhibit, or recording.
- If AI touched a brief, memo, discovery response, client communication, investigation plan, or negotiation document, the responsible attorney remains responsible for the final product.
- If an error reaches a filing, correct it promptly and candidly. *Lnu v. Blanche* shows that attempts to minimize or obscure the source of hallucinated authority can compound the problem.
- Do not bill a client as though a human spent hours doing work that an automated tool produced quickly. Follow applicable fee rules, engagement terms, and disclosure duties.

California SB 574: enacted requirements for generative AI use

SB 574 was chaptered as Chapter 858 and approved by the Governor on September 30, 2026. It takes effect January 1, 2027. For attorneys using generative AI to assist in the practice of law, new Business and Professions Code section 6068.1 and amended Code of Civil Procedure section 128.7 create concrete requirements that should be built into firm workflows now.

- No delegation of the practice of law to generative AI.
- Do not input confidential, personal identifying, or other nonpublic information into a generative AI system unless access to that input is restricted to the attorney and persons the attorney authorizes under confidentiality obligations.
- Take reasonable steps to verify AI outputs, including case and statutory citations, and correct erroneous or hallucinated material that the attorney uses.
- Disclose generative AI use to the court for documents submitted to the court. For public content created using generative AI, consider whether disclosure is appropriate.
- For court filings, an attorney responsible for submission must personally verify citations, including citations supplied by generative AI.
- For arbitrators, SB 574 separately prohibits delegating any part of the decisionmaking process to a generative AI tool and prohibits reliance on AI-generated information outside the record without appropriate disclosure to the parties beforehand and, as far as practical, an opportunity for the parties to comment on its use.

California implementation point

The firm's commercial-workspace rule, source-verification rule, and matter authorization process are directionally consistent with the statute, but they should now be supplemented by a California filing-disclosure step and an express personal-citation-verification step. Client consent is not a substitute for those statutory duties.

7. Build AI into protective orders and discovery protocols

If a case involves confidential discovery, ask about AI before production starts, not after someone has uploaded the entire file to a consumer chatbot and called it workflow optimization.

- Consider language restricting public or open AI tools for discovery materials.
- Define whether confidential material may be used with approved legal AI systems, enterprise systems, local tools, or not at all.
- Address retention, training, deletion, subcontractor access, audit logs, connectors, and who may access prompts and outputs.
- Make sure vendors, experts, investigators, contractors, and staff are covered by the same case-specific restrictions.
- Treat AI use as part of discovery security and preservation, not as a separate technology hobby.

Part 2: Intake questions attorneys should ask clients

These questions belong in modern intake, especially in criminal defense, family law, employment, civil litigation, protective-order matters, and anything involving screenshots, phone extractions, medical records, discovery, or sensitive facts. The companion signable form contains the full acknowledgment, attorney-use disclosure, informed-consent choices, and California workflow checks.

- ☐ Have you used ChatGPT, Claude, Gemini, Copilot, Perplexity, Meta AI, Grok, or another AI tool to discuss this case, including AI built into a phone, browser, search engine, email, document program, or social media app?
- ☐ Did you paste or upload police reports, discovery, body-camera material, jail calls, screenshots, texts, emails, photos, medical or financial records, school records, court documents, or restraining-order paperwork?
- ☐ Did you ask AI whether you could win, what you should say, what defense or claim you might have, or how to explain what happened?
- ☐ Did you use AI to draft a statement, timeline, declaration, email, text message, apology, complaint, or response related to the case?
- ☐ What tool and account did you use: personal, work, school, free, paid, or someone else's account?
- ☐ Did you put anything your attorney told you, or anything from attorney-client conversations, into an AI tool?
- ☐ After you told anyone else about the case, did that person, or anyone acting for you, use AI to discuss the case or enter your information or case materials, with or without your knowledge or permission?
- ☐ Did you delete any AI chats or files?
- ☐ Did you share any AI chat, link, screenshot, summary, or output with anyone else?
- ☐ Did the AI tool tell you to do or not do anything related to the case?
- ☐ Do you still have access to the chats, outputs, uploaded files, or account history?
- ☐ Did any AI use, by you or someone else, happen after you hired or first spoke with your attorney?

Attorney note on third-party AI use

The new question is deliberately broader than “did you use AI?” A parent, spouse, family member, friend, advocate, investigator, or helper may have put the client's facts or documents into an AI system after the client described the case. If the answer is yes or uncertain, identify who used it, which tool and account were used, what information or files were entered, and whether the chat or account still exists. Do not tell the client to ask the other person to delete or alter anything before counsel decides how preservation should be handled.

Collection caution

If AI use is reported, decide the collection method before instructing the client. Counsel may collect directly, direct a controlled export or screenshot, or instruct the client not to access the account further until review. An enthusiastic “I exported everything!” can create its own mess.

Part 3: Scripts attorneys can use with clients**Client script for intake or first meeting**

Please do not upload, paste, summarize, or discuss anything about your case in ChatGPT, Claude, Gemini, Copilot, Perplexity, Meta AI, Grok, or any other AI tool unless I specifically approve it first.

That includes police reports, discovery, screenshots, texts, emails, medical records, court documents, timelines, facts about what happened, witness names, strategy questions, photos, videos, audio, or anything you think might help explain the case. It especially includes anything I have told you or anything from our conversations.

Do not ask a family member, friend, advocate, investigator, or anyone else to put the case into AI for you. If someone else has already done that, or you think they may have, tell me. Do not ask them to delete or change anything until I have reviewed the issue.

AI tools are not your lawyer. They are not automatically confidential. They do not create attorney-client privilege. Deleting a chat may not mean the information is gone.

If you have already used an AI tool for anything related to this case, tell me now so I can evaluate the issue. You are not in trouble. Early is fixable.

Do not forward, screenshot, post, or share any AI chat or output about the case with anyone unless I tell you to, and do not delete anything unless I tell you to.

If you want help organizing your thoughts, send the information to me instead. We can decide what is useful, what is protected, and what should not be placed into a third-party system.

Shorter version for email or engagement letter

Do not use any AI tool to discuss, summarize, upload, analyze, or draft anything about your case unless we approve it first. This includes AI features built into phones, browsers, search engines, email, document programs, and social media apps. Do not ask another person to use AI on the case for you. AI tools are not your lawyer and are not automatically confidential, and putting our communications or advice into them may put protection of those communications at risk. Do not share, forward, alter, or delete AI chats or outputs about your case without our approval. If you, or anyone you told about the case, has already used AI about it, tell us immediately.

Script for disclosing the firm's own approved AI use

Separately from the rules about your own AI use, here is how this office uses AI. We use specifically approved, paid commercial AI systems under written terms and configured access controls. We review what information is appropriate to put into the system, restrict access, review integrations and sharing, and follow the matter's preservation and deletion requirements. I remain responsible for every legal judgment and final work product, and I verify AI-assisted work against the underlying sources. No technology, contract, or setting can guarantee how a court will rule on privilege or work-product questions. Before client-identifying or confidential information is used in an approved AI system, we will discuss the proposed use and document your authorization or limits. California law also imposes separate duties on attorneys using generative AI, including verification, confidentiality safeguards, and, beginning January 1, 2027, disclosure to the court when generative AI is used for documents submitted to the court.

Part 4: Small client handout

Before You Use AI About Your Case

A quick client warning from your legal team

AI tools can be useful for many things. Your case is different. Please do not upload, paste, summarize, or discuss case information in ChatGPT, Claude, Gemini, Copilot, Perplexity, Meta AI, Grok, or another AI tool unless your attorney specifically approves it first. This includes AI built into your phone, browser, search engine, email, document program, or social media app.

Do not put these into AI tools

- Police reports, discovery, body-camera footage, jail calls, phone extractions, screenshots, texts, emails, photos, videos, or audio.
- Medical, mental health, school, financial, employment, immigration, juvenile, or family-law records.
- Court papers, timelines, witness names, strategy questions, private facts, or anything you think might help explain what happened.
- Anything your attorney has told you, or anything from your conversations with your attorney.
- Draft declarations, apologies, statements, complaints, text messages, or emails related to the case.

Do not outsource the problem to someone else

Do not ask a parent, spouse, family member, friend, advocate, investigator, or other helper to put the case into AI for you. If someone else already did, or you think they may have, tell your attorney. Do not ask that person to delete, edit, or close the chat or account.

Why this matters

- An AI tool is not your attorney.
- A chat with AI is not automatically privileged or confidential.
- Deleting a chat may not mean the information is truly gone.
- Prompts, uploads, outputs, or account history can become relevant to discovery, subpoenas, or evidentiary disputes.
- AI can sound confident and still be wrong, because the machine has no natural shame mechanism.

What to do instead

- Send case information to your attorney, not to a personal AI tool.
- If you want help organizing your thoughts, ask your attorney how to do that safely.
- If you or someone else already used AI about the case, tell your attorney now. Early is fixable. Surprise is expensive.
- Do not forward, screenshot, post, share, alter, or delete AI chats or outputs about the case unless your attorney tells you what to do.

What about your legal team's AI use?

A law office may use an approved commercial AI system with contractual, technical, and matter-specific safeguards. That is different from a personal AI app. Your attorney should explain the proposed use, the system and safeguards, and any authorization or limits that apply to your matter.

Bottom line: If it involves your case, ask your attorney before putting it into AI. AI may be helpful. It is not your lawyer, your diary, or a privilege force field.

Part 5: Sample office policy language

This is starter language. It should be reviewed by the attorney responsible for the firm's professional-responsibility and information-security obligations.

Internal AI use policy starter language

No attorney, staff member, contractor, expert, investigator, or vendor may upload, paste, summarize, or discuss client information, case facts, discovery, confidential records, protective-order material, privileged communications, work product, sealed records, juvenile records, medical records, or attorney strategy in any AI tool unless the responsible attorney has approved the specific tool, workspace, settings, and use. Approved use must account for confidentiality, privilege, work product, protective orders, client instructions, retention, training, access controls, integrations, sharing, source verification, and any applicable statutory duties. AI output must be reviewed by the responsible attorney before it is used in client advice, filings, discovery responses, investigation, negotiation, or court work. Unapproved public AI tools, personal AI accounts, browser extensions, third-party GPTs, agents, connectors, or plug-ins may not be used with case materials.

Client work in AI occurs only in the firm's configured commercial workspaces as documented in the firm's setup records. Client-identifying information and raw discovery require matter-specific review and the firm's documented authorization process before upload. Protective orders, sealing restrictions, regulated data, access controls, vendor terms, and retention or preservation requirements must be checked before use.

California SB 574 filing and verification addendum

- No attorney or staff member may delegate legal judgment or the practice of law to a generative AI system.
- Before confidential, personal identifying, or other nonpublic information is entered into generative AI, the responsible attorney must confirm that access to the input is restricted consistent with Business and Professions Code section 6068.1(a)(3)(A).
- AI-assisted outputs must be verified, and any error or hallucination in material the attorney uses must be corrected.
- For every document submitted to a California court on or after January 1, 2027, the responsible attorney must determine whether generative AI was used and complete the disclosure required by section 6068.1(a)(3)(C).
- Before a paper containing citations is filed, an attorney responsible for submission must personally verify each citation as required by amended Code of Civil Procedure section 128.7(b)(2). The amended text also makes section 128.7(d)(1)'s represented-party monetary-sanctions limitation refer to paragraph (2) of subdivision (b), which is now the citation-verification paragraph. Check for later corrective legislation or judicial construction before relying on that cross-reference.
- The file should contain enough documentation to show which attorney completed the verification and court-disclosure step.

Part 6: Current examples attorneys should know

The landscape now includes published appellate discipline decisions in addition to trial-court privilege and discovery rulings. The examples below are for issue spotting, not as a substitute for checking the current law in the controlling jurisdiction.

United States v. Heppner, S.D.N.Y. No. 25-cr-00503-JSR

Heppner is the client-warning example. The defendant used the consumer version of Claude on his own initiative to prepare materials about defense strategy and later sent the outputs to counsel. The district court rejected attorney-client privilege and work-product claims on the facts before it, including the lack of counsel direction and the platform confidentiality problem. The opinion also raised the possibility that putting information learned from counsel into a consumer AI system could affect protection of the underlying attorney communications. The practical lesson is narrower than “AI is always discoverable”: do not let clients assume a consumer chatbot is confidential legal space, and ask about AI use at intake.

California criminal work-product caution. Heppner's work-product analysis turned in part on the fact that the defendant created the AI materials on his own initiative rather than at counsel's direction. Counsel direction therefore may matter, but it is not a magic privilege switch. In California criminal discovery, Penal Code section 1054.6 incorporates only the core written work product described in Code of Civil Procedure section 2018.030(a): a writing that reflects an attorney's impressions, conclusions, opinions, or legal research or theories. The California Supreme Court reaffirmed that limitation in *People v. Nadey* (2024) 16 Cal.5th 102. Accordingly, counsel-directed AI use may strengthen a work-product argument when the prompts or outputs actually reflect counsel's mental processes, but it does not automatically protect every prompt, source document, factual input, upload, or AI-generated output. Separate attorney-client privilege, constitutional, protective-order, and other statutory protections still require their own analysis. This is a California criminal-discovery point; California civil discovery separately provides qualified protection for other attorney work product under Code of Civil Procedure section 2018.030(b).

Lnu v. Blanche, 177 F.4th 1014 (9th Cir. 2026)

The Ninth Circuit disciplined two California attorneys after briefs contained nonexistent cases, misattributed quotations, and serious misstatements of real authority. The court emphasized that it was not sanctioning the mere use of generative AI. The violations arose from signing and filing inaccurate papers and from the later failures of candor about the source of the errors. Each attorney was sanctioned \$2,500 and suspended from Ninth Circuit practice for six months. The firm's attorneys were also ordered for two years to disclose in future filings whether generative AI was used, identify the tool, and certify personal review of the filing and its citations and quotations. The Ninth Circuit also ordered the clerk to serve the discipline order on the State Bar of California and any other applicable licensing authorities for further proceedings as appropriate. The case is a direct warning that verification and candid correction cannot be delegated.

Del Biaggio v. Bansen (2026) 121 Cal.App.5th 831

The California Court of Appeal sanctioned counsel \$1,500 for an opening brief containing misstatements and fictional quotations. Counsel admitted using generative AI for additions to the brief and described a plan for a paralegal to check citations. The court concluded that the protocol was insufficient to satisfy counsel's own obligation to personally review the principal legal authority on which the argument relied. The opinion was ordered forwarded to the State Bar, and counsel was ordered to provide a copy to the client. This is now a California published appellate example of the personal-verification problem.

Warner, Morgan, and Tate: work product can survive, but inputs still matter

Warner v. Gilbarco, Inc., 2026 WL 373043 (E.D. Mich. Feb. 10, 2026), Morgan v. V2X, Inc., 2026 WL 864223 (D. Colo. Mar. 30, 2026), and Tate Group Automotive, LLC v. Legacy Automotive Capital, LLC, No. 25-BC11B-0020 (Tex. Bus. Ct. June 3, 2026) are useful counterweights to the idea that every AI interaction destroys work product. On the facts of those matters, courts recognized protection for AI-assisted work in some circumstances, while still requiring disclosure or accounting concerning the platform or materials supplied to it. The takeaway remains documentation, counsel direction, tool review, and knowing what went into the system.

Jeffries v. Harcros Chemicals: protective orders can reach AI use

Jeffries v. Harcros Chemicals, Inc., No. 25-2569-KHV-ADM, 2026 WL 820218 (D. Kan. Mar. 25, 2026), is the protective-order warning. The court barred open, publicly accessible AI tools for discovery materials and left room for closed, secured environments. The stamp on the document does not answer the upload question. The order does.

The New York Times v. OpenAI: the delete-button lesson, with an important update

The 2025 preservation order is still a useful reminder that ordinary product deletion settings can be displaced by litigation holds. But it should no longer be described as a continuing blanket retention requirement. OpenAI states that its obligations under the earlier order ended September 26, 2025 and that it returned to standard retention practices for new data, while a limited historical April-September 2025 data set remains preserved because of the continuing dispute. The lesson survives, but the current status matters: a user interface deletion control is not the same thing as a guarantee that no legal or security retention exception applies.

California: statute enacted, guidance current, rule amendments still proposed

California now has three distinct layers that should not be collapsed into one. First, SB 574 is enacted law and is scheduled to take effect January 1, 2027. Second, the State Bar's updated 2026 Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law remains the current guidance and addresses agentic AI. Third, COPRAC's proposed amendments to Rules of Professional Conduct 1.1, 1.4, 1.6, 3.3, 5.1, and 5.3 remain proposals on the State Bar's public-comment site. COPRAC approved modified proposals on June 12, 2026 for a second 45-day public-comment period, which closed August 6, 2026. As of October 2, 2026, the State Bar page

continues to identify them as proposed amendments rather than adopted Rules of Professional Conduct.

Part 7: Sources reviewed

Official and current sources should be checked again before finalizing firm policy or relying on a requirement in a specific matter. Key sources used for this October 2026 update include:

[California Legislative Information, SB 574, Chapter 858 \(2026\)](#). Enacted text, including Business and Professions Code section 6068.1 and amended Code of Civil Procedure section 128.7.

[Governor of California, September 30, 2026 AI legislation announcement](#). Confirms SB 574 was signed.

[State Bar of California, 2026 Practical Guidance for the Use of Generative AI in the Practice of Law](#). Updated guidance approved in May 2026, including agentic AI.

[State Bar of California, Proposed AI-related Rules of Professional Conduct amendments](#). Modified proposals approved June 12, 2026; second public-comment deadline August 6, 2026.

[Lnu v. Blanche, Ninth Circuit, June 3, 2026](#). Published appellate discipline order involving AI-related hallucinations, failures of verification and candor, and referral to the State Bar of California.

[Del Biaggio v. Bansen \(2026\) 121 Cal.App.5th 831, A174647](#). Published California appellate opinion imposing sanctions for fictional quotations and inadequate personal verification; the opinion was ordered forwarded to the State Bar.

United States v. Heppner, No. 25-cr-00503-JSR, Dkt. No. 27 (S.D.N.Y. Feb. 17, 2026). Consumer Claude exchanges were not protected on the facts presented, where the defendant acted on his own initiative and not at counsel's direction.

California Penal Code section 1054.6; Code of Civil Procedure section 2018.030(a); People v. Nadey (2024) 16 Cal.5th 102. California criminal discovery limits statutory work-product protection to core written work product reflecting an attorney's impressions, conclusions, opinions, or legal research or theories.

[ABA Formal Opinion 512, Generative Artificial Intelligence Tools \(July 29, 2024\)](#). National ethics baseline addressing competence, confidentiality, communication, supervision, candor, and fees.

[22 NYCRR Part 161, Use of Artificial Intelligence Technology](#). New York court rule effective June 1, 2026.

[OpenAI, response to New York Times data demands](#). Updated status of the 2025 preservation order and standard deletion practices.

[OpenAI, enterprise privacy](#). Business-data training and security commitments.

[Anthropic Privacy Center, commercial model-training policy](#). Claude for Work and API training policy.

[Google Workspace with Gemini privacy information](#). Workspace data protections and model-training restrictions.

[Microsoft, data, privacy, and security for Microsoft 365 Copilot](#). Commercial tenant data protections and retention context.

Existing case examples retained from the prior toolkit: United States v. Heppner, Warner v. Gilbarco, Morgan v. V2X, Tate Group Automotive v. Legacy Automotive Capital, and Jeffries v.

Harcros Chemicals. Attorneys should retrieve the current docket or opinion before relying on any case-specific proposition.

Final line from me

AI can help legal work. It should not be allowed to launder uncertainty into confidence, or confidentiality problems into a neatly formatted paragraph.



VdVLaw
Solutions Beyond Boundaries

Client AI Use Intake Form and Acknowledgment with Attorney AI Consent For Attorneys

Prepared by Robert van der Vijver, Case Cartographer, VdVLaw
VdVLaw | www.vdvlaw.net
Current as of October 2, 2026

Attorney-directed legal support. Not a law firm. No legal advice.

Template prepared by Robert van der Vijver, Case Cartographer, VdVLaw (vdrvlaw.net), an attorney-directed legal support service. VdVLaw is not a law firm and does not provide legal advice. The supervising attorney is responsible for reviewing, adapting, and approving this form before use.

Client AI Use Intake Form and Acknowledgment

Attorney template. Complete with the client at intake or first meeting. Retain the signed original in the client file.

Client name: _____

Matter / case: _____

Attorney / firm: _____

Date completed: _____

Completed by (attorney / staff): _____

Why we ask these questions

AI tools like ChatGPT, Claude, Gemini, and Copilot are useful for many things, and many people use them to think through stressful situations. That is normal, and you are not in trouble. But what goes into those tools may not be private, may not be protected, and may become part of the case, including through discovery or a subpoena. We ask these questions so we can find and handle any issue early. Please answer completely and honestly. AI tools include features built into search engines, phones, browsers, email, document programs, and social media apps, not only standalone chatbot apps.

Section 1: AI use questions

1. Have you used ChatGPT, Claude, Gemini, Copilot, Perplexity, Meta AI, Grok, or any other AI tool to discuss anything about this case or situation? This includes AI built into your phone, browser, search engine, email, document program, or social media app.

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

2. Which tool or tools did you use, and what kind of account was it: personal, work, school, free, paid, or someone else's account? Account email or username used, if known:

Details:

3. Did you paste or upload any police reports, discovery, body camera transcripts, jail calls, screenshots, text messages, emails, photos, medical records, financial records, school records, court documents, or restraining-order paperwork?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

4. Did you ask the AI tool whether you could win, what you should say, what defense or claim you might have, or how to explain what happened?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

5. Did you use AI to draft a statement, timeline, declaration, email, text message, apology, complaint, or response related to the case?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

6. Did you upload files, photos, audio, video, or screenshots?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

7. Did you put anything your attorney has told you, or anything from your conversations with your attorney, into an AI tool?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

8. After you told anyone else about this case or situation, did that person, or anyone else acting for you, use ChatGPT, Claude, Gemini, Copilot, or another AI tool to discuss the case or enter any of your information or case materials, with or without your knowledge or permission?

☐ Yes ☐ No ☐ Not sure

If yes or not sure: Who used it? Which AI tool and account, if known? What information or materials were entered? Is the chat, output, or account history still accessible?

9. Did you delete any AI chats or files after using the tool?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

10. Did you share any AI chat, link, screenshot, summary, or output with anyone else?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, with whom and what was shared?

11. Did the AI tool tell you to do, or not do, anything related to the case?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

12. Do you still have access to the chats, outputs, uploaded files, or account history?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

13. Did any AI use, by you or someone else, happen after you hired or first spoke with your attorney?

☐ Yes ☐ No ☐ Not sure

If yes or not sure, details:

Important preservation instruction

Do not delete, edit, or close any AI account, chat, or file discussed above unless your attorney tells you to. If someone else may have used AI about your case, do not ask that person to delete, alter, or close anything. Tell your attorney and let the attorney decide how the material should be reviewed or preserved.

Section 2: Instruction to client regarding AI tools

The attorney or staff member reviews this instruction with the client before signing.

Please do not upload, paste, summarize, or discuss anything about your case in ChatGPT, Claude, Gemini, Copilot, Perplexity, Meta AI, Grok, or any other AI tool unless I specifically approve it first.

That includes police reports, discovery, screenshots, texts, emails, medical records, court documents, timelines, facts about what happened, witness names, strategy questions, photos, videos, audio, or anything you think might help explain the case. It especially includes anything I have told you or anything from our conversations.

Do not ask a parent, spouse, family member, friend, advocate, investigator, or other helper to put this case into AI for you. If anyone else already has, or you think they may have, tell me. Do not ask that person to delete or change anything.

AI tools are not your lawyer. They are not automatically confidential. They do not create attorney-client privilege. Deleting a chat may not mean the information is gone.

If you have already used an AI tool for anything related to this case, tell me now so I can evaluate the issue. You are not in trouble. Early is fixable.

Do not forward, screenshot, post, share, alter, or delete any AI chat or output about the case unless I tell you to.

If you want help organizing your thoughts, send the information to me instead. We can decide what is useful, what is protected, and what should not be placed into a third-party system.

Section 3: Attorney Use of Approved AI and Client Informed Consent

The attorney should personally review this section with the client before obtaining authorization. ABA Formal Opinion 512 and applicable state rules govern the informed-consent analysis. For California matters, the attorney should also account for Business and Professions Code section 6068.1 and amended Code of Civil Procedure section 128.7, effective January 1, 2027.

Important distinction

The client is instructed not to place case information into personal or consumer AI tools. This section concerns the attorney's separate use of a specifically approved commercial AI workspace that the attorney has reviewed and configured for law-firm use.

Plain-language summary for the client

In plain terms: the firm uses a specifically identified, paid commercial AI platform under written terms and configured safeguards. The attorney reviews the actual workspace, access, integrations, sharing, retention, and data-handling terms. The provider may process information to deliver the service. These safeguards are intended to protect confidentiality and reduce risk, but no technology, contract, setting, consent form, or statute can guarantee how a court or other authority will decide privilege or work-product questions.

Your attorney remains your lawyer. AI output is a starting point, never the finished product. The attorney retains legal judgment, reviews and verifies AI-assisted work, and remains responsible for the representation, the advice you receive, and everything filed or sent on your behalf.

Approved system and proposed use

AI platform and commercial plan: _____

Firm workspace / tenant name: _____

Expected uses in this matter: _____

Types of information that may be processed: _____

Retention or deletion terms explained: _____

Expected effect on fees or costs: _____

California court-filing disclosure workflow, if applicable: _____

What the attorney has explained

- **Purpose and benefit.** The approved system may assist with organizing and analyzing discovery, summarizing records, comparing documents, preparing timelines and outlines, drafting, and issue spotting. The attorney remains responsible for every legal judgment and final work product.
- **Confidentiality safeguards.** The firm uses a commercial work account with contractual and technical safeguards. For California matters subject to section 6068.1, the attorney must confirm that access to confidential, personal identifying, and other nonpublic information entered into generative AI is restricted to the attorney and persons the attorney authorizes under confidentiality obligations.
- **Privilege and work product.** The attorney intends to treat information placed in the approved system as confidential and, where applicable, protected by attorney-client privilege or work product. This authorization is not intended to waive those protections. Privilege and work product remain legal determinations under applicable law.
- **Material risks.** The provider must process information to provide the service and may retain data under contract terms, security procedures, product settings, legal obligations, or preservation requirements. AI output can be incomplete or wrong. Account compromise, accidental sharing, unapproved integrations, and future product changes create additional risk.
- **Reasonably available alternatives.** The attorney may perform work manually, use another approved technology, exclude identifying information, limit document categories, or decline AI use for the matter. These alternatives may affect speed, cost, or usefulness.
- **Fees and costs.** The attorney has explained any expected effect of the proposed AI use on legal fees, expenses, or technology charges, if applicable.
- **Verification and filing duties.** The attorney will verify AI-assisted legal and factual material before use. For California court papers subject to SB 574, the responsible attorney must personally verify citations and complete the required court disclosure if generative AI was used.
- **Limits on use.** The attorney will not knowingly submit information barred by a protective order, sealing order, law, regulation, client instruction, or vendor contract. Regulated information will be used only when required agreements and safeguards are in place.

Client authorization

After discussing the proposed use, benefits, material risks, safeguards, and alternatives with my attorney, I choose one:

☐ I authorize my attorney and supervised law-firm personnel to use the approved commercial AI system described above for this matter, including processing client names, identifying information, confidential information, privileged communications, attorney work product, and other matter information reasonably necessary for the representation, excluding raw discovery and case files unless separately authorized below, subject to the limits stated in this form. Client initials: ____

☐ I authorize the attorney's use of the approved system only with these limits: _____ Client initials: ____

☐ I do not authorize client-identifying or confidential information to be placed into an AI system. The attorney may use AI only with public, fictional, or sufficiently anonymized information unless I later provide informed consent. Client initials: ____

Separate authorization for raw discovery and case files: ☐ I authorize the attorney to upload and analyze discovery and case files containing names or identifying information when the attorney determines that doing so is reasonably necessary and permitted by protective orders, sealing requirements, law, and the platform's commercial terms. If this box is not selected, the general authorization above does not permit raw discovery or case-file uploads. Client initials: ____

I understand that I may revoke this authorization prospectively by notifying my attorney in writing. Revocation cannot undo processing that occurred before the attorney received the revocation or retention required by contract, firm policy, law, court order, security procedure, or preservation obligation.

Attorney certification

☐ I personally discussed the proposed AI use with the client, including the specific system, information that may be processed, expected benefit, material risks, reasonably available alternatives, and the limits selected above.

☐ I reviewed the actual commercial workspace, access controls, sharing, connectors or integrations, retention settings, and applicable vendor terms for the intended use.

☐ For California matters subject to Business and Professions Code section 6068.1, I confirmed the restricted-access requirement for confidential, personal identifying, and other nonpublic information before use.

☐ I understand that I may not delegate the practice of law or legal judgment to generative AI and that I remain responsible for verification and correction of AI-assisted material.

☐ For California court submissions on or after January 1, 2027, I will complete the generative-AI disclosure required by section 6068.1(a)(3)(C) when applicable and will personally verify citations as required by Code of Civil Procedure section 128.7(b)(2).

Attorney initials: _____ Date: _____ Jurisdiction-specific authority reviewed:

Authority note: ABA Formal Opinion 512 explains informed-consent principles under the Model Rules. For California matters, counsel should also review California Rules of Professional Conduct 1.0.1(e), 1.4, 1.6, 3.3, 5.1, and 5.3; Business and Professions Code sections 6068(e)(1) and 6068.1; Code of Civil Procedure section 128.7; and the State Bar of California's current Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law. The supervising attorney must adapt this section to the controlling jurisdiction, actual system, and circumstances of the matter.

Section 4: Client acknowledgment and signatures

- ☐ ___ Section 3 was reviewed with me by the attorney. I received information about the proposed AI use, material risks, safeguards, and reasonably available alternatives, and my authorization choice is recorded there.
- ☐ ___ I understand that the attorney's authorized use of an approved commercial AI workspace does not authorize me to use a personal or consumer AI tool for this case.
- ☐ ___ My answers in Section 1 are true and complete to the best of my knowledge, including anything I know or suspect about AI use by another person after I discussed the case with them.
- ☐ ___ The instruction in Section 2 was reviewed with me, I had the chance to ask questions, and I understand it.
- ☐ ___ I understand that AI tools are not my lawyer, that my communications with them are not automatically confidential or privileged, and that deleting a chat may not remove the information.
- ☐ ___ I agree not to use any AI tool, or ask another person to use an AI tool, to discuss, summarize, upload, analyze, or draft anything about my case unless my attorney approves it first.
- ☐ ___ I agree to tell my attorney promptly if I remember, discover, or suspect past AI use related to my case by me or anyone else, or if any new AI use happens.
- ☐ ___ I agree not to delete, edit, close, or ask someone else to delete, edit, or close any AI account, chat, or file related to my case unless my attorney instructs me to.

Client signature

Date

Client printed name

Attorney signature (Section 3 discussed with client)

Date

Attorney use only: follow-up

- ☐ ___ No AI use reported; instruction given and acknowledged.
- ☐ ___ AI use reported by client; preservation instruction given.
- ☐ ___ Possible AI use by another person identified; who, tool/account, information entered, and current access status documented.
- ☐ ___ Client instructed not to ask any third party to delete, alter, or close AI chats, files, or accounts before counsel decides preservation steps.
- ☐ ___ Chats, outputs, uploaded files, or account records to be collected and reviewed.
- ☐ ___ Collection method selected: counsel collects directly, controlled export or screenshot, or instruct client not to access the account further until reviewed.
- ☐ ___ Privilege, work product, protective order, confidentiality, disclosure, and preservation issues evaluated.
- ☐ ___ Section 3 informed-consent discussion completed by the attorney; client authorization choice documented.
- ☐ ___ Expected effect on fees, expenses, or technology charges explained and documented, if applicable.
- ☐ ___ Approved commercial platform, workspace, intended uses, data categories, access restrictions, connectors, sharing, and retention terms identified.
- ☐ ___ Any limits or refusal of authorization flagged in the matter-management system and communicated to staff.
- ☐ ___ Protective orders, regulated data, sealing requirements, and vendor-contract requirements reviewed before any raw discovery upload.
- ☐ ___ California SB 574 check, if applicable: no delegation; restricted access confirmed; AI output verified and corrected; court disclosure completed for AI-assisted court documents on or after January 1, 2027; citations personally verified under Code of Civil Procedure section 128.7.
- ☐ ___ Calendar follow-up. Date: _____

Notes:

Template customization note: Replace the header and footer with firm information and add jurisdiction-specific language before use. The attorney must review the actual platform contract,

privacy terms, access restrictions, retention settings, protective orders, statutory requirements, and applicable professional-responsibility rules. This form documents a discussion and authorization; it does not itself guarantee confidentiality, privilege, work-product protection, or compliance.