SUPPLEMENTAL PUBLIC COMMENT ON COPRAC'S JUNE 2026 MODIFIED AI RULE PROPOSALS

From Professional Accountability to Workable Oversight

Submitted to:
The State Bar of California
Standing Committee on Professional Responsibility and Conduct (COPRAC)

Submitted by:
Jonathan Z. Morris, Esq. (SBN 345401)
Gordon Rees Scully Mansukhani, LLP
633 West Fifth Street, 52nd Floor
Los Angeles, California 90071
jzmorris@grsm.com | (213) 576-5091
Submitted in Individual Capacity

August 6, 2026

Affiliation and contact information are provided for identification purposes only. The views expressed are solely those of the author and do not necessarily reflect the views of Gordon Rees Scully Mansukhani, LLP, its attorneys, or its clients.

Dear Members of COPRAC:

I appreciate the opportunity to provide supplemental comments on the proposed amendments to Rules 1.1, 1.4, 1.6, 3.3, 5.1, and 5.3 of the California Rules of Professional Conduct addressing the use of artificial intelligence. I submitted comments during the initial public-comment period supporting COPRAC's effort to apply existing professional duties to AI-assisted practice and recommending targeted clarification of the proposed communication and confidentiality standards. Following the close of the initial public-comment period on May 4, 2026, COPRAC revised the proposed amendments in response to the comments received. At its June 12, 2026 meeting, COPRAC approved the modified rule proposals for a second 45-day public-comment period ending August 6, 2026 (the “June 2026 Modified Proposal”).

The June 2026 Modified Proposal materially improves the framework. It correctly treats artificial intelligence as a form of technology governed by existing duties rather than as a separate category of professional conduct. Furthermore, the revised confidentiality comment appropriately focuses on the risk that protected information may be accessed, retained, or used inconsistently with the lawyer's duty. The proposal also properly reinforces the nondelegable duties of competence, candor, communication, and supervision.

The remaining questions are now more operational. The central issue is no longer whether a lawyer remains responsible when using AI. The lawyer plainly does. The unresolved issue is what responsible review, supervision, and approval require when technology does more than generate a single draft. That question becomes more difficult when a system retrieves information, makes intermediate classifications, recommends a course of action, interacts with connected systems, or completes portions of a multi-stage workflow with limited additional prompting.

The future ethical question will not be merely, “Did the lawyer use AI?” It will be whether the lawyer can explain, defend, and stand behind every material decision, factual proposition, legal conclusion, recommendation, and external action that AI helped produce.

This supplemental comment therefore supports the June 2026 Modified Proposal while recommending five focused refinements designed to make the framework administrable, technologically realistic, and coherent with emerging California law. In summary, COPRAC should: (1) replace Rule 1.1's unqualified "all aspects" formulation with a materiality-focused standard; (2) expressly adopt risk-proportionate review and verification; (3) add an operative disclosure threshold under Rule 1.4; (4) provide a control and approval framework for agentic and integrated systems; and (5) coordinate the Rules with the overlapping statutory duties contemplated by Senate Bill 574 and other applicable law.

I. THE MODIFIED PROPOSAL GETS THE CORE PRINCIPLE RIGHT

The June 2026 Modified Proposal rests on the correct premise: artificial intelligence changes how legal work may be performed, but it does not change who is professionally responsible for that work. Proposed Comment [2] to Rule 1.1 requires a lawyer using technology, including artificial intelligence, to exercise professional judgment over "all aspects" of that use, including inputs and outputs. Proposed Rule 3.3 Comment [3] similarly requires verification of cited authority before submission to a tribunal. Proposed amendments to Rules 5.1 and 5.3 reinforce managerial and supervisory responsibility for firm policies, lawyer conduct, and nonlawyer use of technology.

This approach is preferable to a standalone AI rule built around a transient definition of a particular technology. The Rules of Professional Conduct should remain durable as products, interfaces, and terminology change. Competence, confidentiality, communication, candor, and supervision are stable obligations. The challenge is to describe how those obligations apply when the technology's role becomes more autonomous, integrated, and consequential.

The 2026 Practical Guidance recognizes that development. It expressly addresses agentic AI - systems that may autonomously perform tasks or workflows without repeated human prompting - and explains that lawyers must understand the nature and extent of a system's autonomy, the circumstances in which it may act, the data sources it may access, and the safeguards necessary to prevent error and ethical violations. The proposed rule comments should more directly translate that practical guidance into an administrable standard.

II. THE DISCUSSION HAS MOVED BEYOND HALLUCINATED CITATIONS

Recent practitioner commentary reflects broad agreement on several foundational points, including that (1) lawyers remain responsible for work performed with technological assistance; (2) legal authority and material facts must be verified; (3) confidential information must be protected; and (4) firms must establish appropriate policies, training, and oversight. The emerging disagreement concerns implementation: what must be reviewed, when must approval occur, and how should oversight scale with a system's autonomy and the consequences of error?

A. Hallucinations are a visible symptom, not the entire problem

Fabricated cases and quotations have understandably driven much of the regulatory response. But an AI system may produce accurate text and still create professional risk. It may omit a dispositive record, rely on data the lawyer was not authorized to disclose, generate a recommendation from incomplete sources, initiate an external communication without meaningful approval, or produce an apparently reasonable answer that the lawyer adopts without understanding the assumptions embedded in it. Accuracy is essential, but accountability extends beyond accuracy.

B. Accountability requires identifiable human judgment

A useful formulation circulating in practitioner discussions is that AI should assist, a lawyer should approve, and the organization should be able to audit consequential use. That concept should not be converted into a universal mandate to preserve every prompt or reconstruct every computational step. It does, however, identify the correct regulatory focus: material legal judgment must remain with an identifiable lawyer, consequential external actions should have a meaningful approval point, and material factual and legal propositions should be traceable to reliable sources.

C. Agentic systems expose the ambiguity of an unqualified "all aspects" standard

A conventional drafting tool receives instructions and returns a draft. An agentic or integrated system may instead determine a sequence of tasks, retrieve information from multiple repositories, classify or prioritize documents, select tools, revise its own approach, and generate a final product after many intermediate operations. Proposed Rule 1.1 Comment [2] requires a lawyer to exercise professional judgment over "all aspects" of technology use, including inputs and outputs. That formulation appropriately recognizes that professional risk may arise before a final output is produced. But it does not identify which aspects are material, what level of review is required, or how the duty applies when the relevant contribution is an intermediate recommendation, an automated action, the selection or omission of sources, or access to connected systems.

Two extreme interpretations should be avoided. Reading "all aspects" literally could imply a duty to reconstruct hidden, technical, or immaterial computational operations that a lawyer cannot meaningfully inspect. Reading it as requiring review only of the final endpoint may be inadequate for a high-risk autonomous workflow. The governing standard should therefore focus on materiality, risk, traceability, and control.

D. AI risk extends beyond the lawyer's own prompts

Practitioner commentary has also identified a client-side risk not squarely addressed by the proposed comments. A client may independently enter a chronology, factual narrative, document, or litigation strategy into a public consumer AI system and later provide the resulting material to counsel. Recent federal decisions discussed in the California legal press have reached different conclusions concerning privilege and work product based on facts such as whether the use was unilateral or attorney-directed, whether the system was public or protected by contractual and technical safeguards, and whether confidentiality existed at the time of the communication. These developments support practical guidance advising lawyers, when reasonably foreseeable, to educate clients not to place matter-related information into public or unapproved AI systems without consulting counsel.

III. SB 574 CREATES A FORESEEABLE STATUTORY OVERLAY

The need for clarity is heightened by Senate Bill 574. As of this submission, the July 2, 2026 amended version remains pending in the Assembly. The bill passed the Senate 39-0, received unanimous votes in the Assembly Judiciary and Privacy and Consumer Protection Committees, and was placed on the Assembly Appropriations suspense file on August 5, 2026. The bill may be amended further and should not be treated as enacted. Its progress nevertheless demonstrates that California is simultaneously developing overlapping professional, statutory, and tribunal-based duties governing AI use.

The current bill would prohibit an attorney from delegating the practice of law to generative AI; restrict entry of confidential, personal identifying, and other nonpublic information into systems lacking appropriate access restrictions; require reasonable steps to verify AI outputs and correct erroneous or hallucinated material; and amend Code of Civil Procedure section 128.7 to require personal reading and verification of cited authorities in court filings. The July 2 text also includes a broad court-disclosure provision for documents created with generative AI, although the legislative process remains ongoing.

Those provisions substantially overlap with the proposed comments, but they use different formulations and may create different triggers. Lawyers may be required to reconcile: professional judgment over “all aspects” of technology use under Rule 1.1; a reasonable-steps verification standard under proposed Business and Professions Code section 6068.1; personal review of citations under Code of Civil Procedure section 128.7; risk-based confidentiality under Rule 1.6; categorical restrictions concerning certain public systems; and any separate tribunal disclosure requirement.

COPRAC need not copy a moving bill into permanent rule comments. It should, however, make the Rules compatible with foreseeable statutory obligations. The final framework should explain that technological assistance is permissible when the lawyer retains and exercises independent professional judgment over material legal decisions; that external statutory, court, and tribunal requirements apply according to their terms; and that compliance with a specific statute does not necessarily satisfy the lawyer's broader professional duties.

IV. FIVE IMPLEMENTATION PROBLEMS AND PROPOSED SOLUTIONS

1. Rule 1.1 should tie the "all aspects" formulation to materiality

Proposed Comment [2] requires a lawyer to exercise professional judgment over "all aspects" of technology use, including inputs and outputs. That breadth recognizes that competence extends beyond final text, but it may be read to encompass technical, hidden, or immaterial operations that no lawyer can meaningfully inspect. The ambiguity becomes more significant when a system's material contribution is a recommendation, classification, source-selection decision, automated communication, or other action rather than a traditional written output.

The comment should therefore focus the duty on material aspects of technology use and expressly include material inputs, outputs, recommendations, decisions, and actions. The goal is not to diminish lawyer responsibility. It is to define that responsibility around the parts of the technology's use that affect the representation and can be meaningfully evaluated. The proposed language appears in Part V.A below.

2. The nature and extent of review should be proportionate to risk

The same review protocol should not govern every use. Formatting assistance, internal brainstorming, a source-linked medical chronology, a strategic recommendation, client advice, and an autonomous external communication present materially different risks. A workable standard should account for the significance of the task, the system's autonomy and access, the traceability of its sources, whether the result will be communicated externally, the reversibility of the action, and the foreseeable consequences of error.

A proportional standard would protect clients without requiring lawyers to reproduce manually every operation performed by a system. It would also provide a principled answer to competing practitioner formulations such as “verify every word” and “review only the final output.” Neither absolute is adequate in every context. The proposed language appears in Part V.A below.

3. Rule 1.4 should identify the operative disclosure threshold

Proposed Rule 1.4 Comment [5] directs lawyers to evaluate their communication obligations concerning technology throughout the representation based on the facts and circumstances, including the novelty of the technology, its risks and benefits, the scope of the representation, and the sophistication of the client. The contextual approach is sound, but the comment does not identify when that evaluation results in an actual duty to communicate.

Without an operative trigger, lawyers may either disclose every routine AI-assisted task out of caution or disclose nothing because the comment does not explain when the listed considerations become significant enough to require communication. The comment should require communication when AI use presents a significant risk to the client’s interests or materially affects the scope, cost, manner, or decision-making process of the representation, while clarifying that routine, lawyer-directed use without such a risk or effect does not independently require disclosure. The proposed language appears in Part V.C below.

4. Rules 5.1 and 5.3 should translate agentic autonomy into a control and approval framework

Rule 5.1 appropriately requires firm policies governing AI, and Rule 5.3 appropriately extends instruction and supervision to nonlawyer use of technology. But the proposed comments regulate people who use systems without describing reasonable controls for systems capable of acting across a workflow. The issue is particularly important where a centrally approved system has continuing access to firm repositories, can initiate tasks, can communicate through connected applications, or can materially affect a representation without a new prompt at each step.

Reasonable measures will vary by firm and use case. They may include defined authorization limits, role-based access, approved data sources, source traceability, human approval before material external communications, filings, advice, or irreversible actions, testing and monitoring, escalation procedures, and periodic review of vendor or system changes. These examples would make Rule 5.1 operational without converting every prompt into a regulated event. The same framework should clarify that individual lawyers may rely reasonably on firm-approved tools and controls absent warning signs, while managerial lawyers or designated governance bodies remain responsible for reasonable vendor diligence, configuration, access controls, training, and monitoring. The proposed language appears in Part V.E below.

5. The Rules should expressly coordinate with SB 574 and other applicable law

The Rules should anticipate overlapping statutory and tribunal requirements without incorporating the current version of SB 574 verbatim. A short harmonization comment can clarify permissible assistance, preserve the lawyer's nondelegable judgment, and explain that specific legal requirements supplement rather than replace professional duties.

The harmonization should address three points: technological assistance is not itself an impermissible delegation where the lawyer retains material legal judgment; statutory, court, client, and tribunal requirements apply according to their terms; and compliance with an external requirement does not necessarily establish compliance with broader professional duties. The proposed language appears in Parts V.B, V.D, and V.E below.

V. PROPOSED REVISIONS BY RULE

The following comparisons separate the current COPRAC proposal from the language recommended in this supplemental comment. Material changes to the current proposed text are shown in bold. Appendix A presents the recommended language in clean form.

A. Rule 1.1, Comment [2] – Material Aspects and Proportional Review

CURRENT PROPOSED TEXT

[2] When using technology, including artificial intelligence, competence requires a lawyer to exercise professional judgment over all aspects of that use, including, but not limited to, the inputs and outputs. When citing legal authority to a tribunal, a lawyer must comply with the duty of candor. See rule 3.3; Business and Professions Code section 6068(d). When citing legal authority to a client or to other parties, the lawyer must verify the accuracy and existence of cited authorities to ensure the lawyer employs only those means consistent with the truth. See Business and Professions Code section 6068(d).

PROPOSED REVISION

[2] When using technology, including artificial intelligence, competence requires a lawyer to exercise professional judgment over **the material aspects of that use, including material inputs and outputs and any material recommendations, decisions, or actions generated or taken by the technology. The nature and extent of review, verification, and supervision should be reasonable under the circumstances and should account for the materiality of the task, the technology's degree of autonomy and access, the traceability and reliability of the information used, whether the result will be communicated, relied upon, or acted upon externally, and the foreseeable consequences of error.** When citing legal authority to a tribunal, a lawyer must comply with the duty of candor. See rule 3.3; Business and Professions Code section 6068(d). When citing legal authority to a client or to other parties, the lawyer must verify the accuracy and existence of cited authorities to ensure the lawyer employs only those means consistent with the truth. See Business and Professions Code section 6068(d).

PURPOSE OF REVISION

Replaces the unqualified "all aspects" formulation with a materiality-focused standard, makes clear that competence extends beyond final text to consequential recommendations, decisions, and actions, and establishes review proportionate to the technology's risk and autonomy.

B. Rule 1.1 - Nondelegation and External-Law Overlay

CURRENT PROPOSED TEXT

No corresponding language.

PROPOSED ADDITION

The use of technology to assist in performing legal tasks does not constitute delegation of the practice of law where the lawyer retains and exercises independent professional judgment over material legal decisions and reviews and approves material work product, recommendations, and actions before they are used or taken in the representation. The duties stated in these rules operate in addition to applicable statutes, rules of court, standing orders, tribunal directives, and other law. Compliance with a specific external requirement does not, by itself, establish compliance with the lawyer's broader duties under these rules.

PURPOSE OF REVISION

Clarifies permissible technological assistance, preserves the lawyer's nondelegable professional judgment, and anticipates overlapping statutory and tribunal requirements.

C. Rule 1.4, Comment [5] - Operative Disclosure Threshold

CURRENT PROPOSED TEXT

[5] A lawyer's duty to keep a client reasonably informed includes evaluation of the lawyer's communication obligations concerning the lawyer's use of technology, including artificial intelligence, throughout the representation based on the facts and circumstances, including the novelty of the technology, risks and benefits associated with the use of the technology, scope of the representation, and sophistication of the client.

PROPOSED REVISION

[5] A lawyer's duty to keep a client reasonably informed includes evaluation of the lawyer's communication obligations concerning the lawyer's use of technology, including artificial intelligence, throughout the representation based on the facts and circumstances, including the novelty of the technology, risks and benefits associated with the use of the technology, scope of the representation, and sophistication of the client. **Communication concerning the use of technology is required when the use presents a significant risk to the client's interests or materially affects the scope, cost, manner, or decision-making process of the representation. Routine, lawyer-directed use that does not present such a risk or material effect does not, standing alone, require separate disclosure under this rule.**

PURPOSE OF REVISION

Adds an operative disclosure trigger to COPRAC's contextual framework and clarifies that routine, lawyer-directed assistance does not independently require disclosure.

D. Rule 3.3, Comment [3] - Verification and Separate Disclosure Requirements

CURRENT PROPOSED TEXT

[3] A lawyer's duty of candor towards the tribunal includes the obligation to verify the accuracy and existence of cited authorities, including ensuring no cited authority is fabricated, misstated, or taken out of context, before submission to a tribunal, including any cited authorities generated or assisted by artificial intelligence or other technological tools.

PROPOSED REVISION

[3] A lawyer's duty of candor towards the tribunal includes the obligation to verify the accuracy and existence of cited authorities, including **taking reasonable steps to ensure** no cited authority is fabricated, misstated, or taken out of context, before submission to a tribunal, including any cited authorities generated or assisted by artificial intelligence or other technological tools. **Nothing in this comment independently requires disclosure of the use of technology to a tribunal unless disclosure is required by applicable law, rule of court, standing order, or tribunal directive. Any separate requirement concerning review, verification, certification, or disclosure must be satisfied according to its terms.**

PURPOSE OF REVISION

Preserves a meaningful verification duty while distinguishing reasonable professional review from an implied guarantee of perfect accuracy and from any separate statutory or tribunal disclosure requirement.

E. Rule 5.1, Comment [1] - Agentic Governance and External Requirements

CURRENT PROPOSED TEXT

[1] Paragraph (a) requires lawyers with managerial authority within a law firm to make reasonable efforts to establish internal policies and procedures designed, for example, to detect and resolve conflicts of interest, identify dates by which actions must be taken in pending matters, account for client funds and property, ensure that inexperienced lawyers are properly supervised, and govern the use of artificial intelligence, in accordance with the Rules of Professional Conduct.

PROPOSED REVISION

[1] Paragraph (a) requires lawyers with managerial authority within a law firm to make reasonable efforts to establish internal policies and procedures designed, for example, to detect and resolve conflicts of interest, identify dates by which actions must be taken in pending matters, account for client funds and property, ensure that inexperienced lawyers are properly supervised, and govern the use of artificial intelligence, in accordance with the Rules of Professional Conduct. **Policies governing the use of artificial intelligence should account for the technology's degree of autonomy, access to protected information and connected systems, ability to act or communicate externally, and capacity to affect substantive legal decisions. Reasonable measures may include authorization limits, access controls, source-traceability requirements, testing and monitoring, escalation procedures, and meaningful lawyer approval before material external communications, filings, advice, or other consequential actions. Reasonable firm policies should also identify and account for applicable statutory, court-imposed, client-imposed, and tribunal-specific**

requirements concerning technological assistance, including requirements governing protected information, verification, approval, certification, and disclosure.

PURPOSE OF REVISION

Makes the general policy obligation operational for integrated and agentic systems and requires firms to account for the growing external-law overlay.

F. Rules 1.6 and 5.3 - No Further Textual Revision Requested

This supplemental comment supports Rule 1.6 as revised. Its risk-based treatment of access, retention, use, and confidentiality is a substantial improvement. It also supports the proposed Rule 5.3 clarification concerning instruction and supervision of nonlawyers using technology. The additional governance measures proposed above can be implemented through Rule 5.1 without further revising Rule 5.3.

VI. RULE-SPECIFIC POSITIONS

PROVISION	POSITION	RECOMMENDED ACTION
Rule 1.1 - Competence	Support if modified	Replace the unqualified "all aspects" formulation with a materiality-focused standard; add risk-proportionate review; clarify nondelegation and the external-law overlay.
Rule 1.4 - Communication	Support if modified	Add an operative disclosure trigger keyed to significant risk or material effect, together with a routine-use clarification.
Rule 1.6 - Confidentiality	Support	Preserve the revised risk-based framework. Address remaining operational issues through practical guidance.
Rule 3.3 - Candor	Support if clarified	Preserve verification; use a reasonable-steps formulation; distinguish separate certification and disclosure requirements.
Rule 5.1 - Managerial/Supervisory Lawyers	Support if modified	Add agentic governance, approval controls, allocation of firm-level responsibility, and tracking of external requirements.
Rule 5.3 - Nonlawyer Assistants	Support	Preserve the proposed instruction and supervision language; rely on Rule 5.1 for firm-level controls.

VII. CONCLUSION

COPRAC has correctly established that a lawyer cannot outsource professional responsibility to technology. The next step is to make that principle workable as AI becomes integrated into the entire lifecycle of a representation rather than confined to a single drafting task.

California should not measure responsible AI use by whether a lawyer can reproduce every computational operation performed by a technological system. It should ask whether the lawyer understood the system's role, imposed controls proportionate to its risk and autonomy, independently evaluated the material result, approved consequential actions, protected client information, complied with applicable external requirements, and remained able to explain, defend, and stand behind the representation.

The five refinements proposed here would preserve the nondelegable nature of legal judgment while providing lawyers and law firms with an administrable standard. They would also position the Rules to remain coherent if SB 574 or related statutory and tribunal requirements become law in a different form. Accordingly, I respectfully request that COPRAC adopt these clarifications before forwarding the proposed amendments for further consideration.

Respectfully submitted,

Jonathan Z. Morris (SBN 345401)
Partner, Gordon Rees Scully Mansukhani, LLP
Submitted in Individual Capacity

APPENDIX A - CLEAN PROPOSED LANGUAGE

Rule 1.1, Comment [2] - Material Aspects and Proportional Review

[2] When using technology, including artificial intelligence, competence requires a lawyer to exercise professional judgment over the material aspects of that use, including material inputs and outputs and any material recommendations, decisions, or actions generated or taken by the technology. The nature and extent of review, verification, and supervision should be reasonable under the circumstances and should account for the materiality of the task, the technology's degree of autonomy and access, the traceability and reliability of the information used, whether the result will be communicated, relied upon, or acted upon externally, and the foreseeable consequences of error. When citing legal authority to a tribunal, a lawyer must comply with the duty of candor. See rule 3.3; Business and Professions Code section 6068(d). When citing legal authority to a client or to other parties, the lawyer must verify the accuracy and existence of cited authorities to ensure the lawyer employs only those means consistent with the truth. See Business and Professions Code section 6068(d).

Rule 1.1 - Nondelegation and External-Law Overlay

The use of technology to assist in performing legal tasks does not constitute delegation of the practice of law where the lawyer retains and exercises independent professional judgment over material legal decisions and reviews and approves material work product, recommendations, and actions before they are used or taken in the representation. The duties stated in these rules operate in addition to applicable statutes, rules of court, standing orders, tribunal directives, and other law. Compliance with a specific external requirement does not, by itself, establish compliance with the lawyer's broader duties under these rules.

Rule 1.4, Comment [5] - Operative Disclosure Threshold

[5] A lawyer's duty to keep a client reasonably informed includes evaluation of the lawyer's communication obligations concerning the lawyer's use of technology, including artificial intelligence, throughout the representation based on the facts and circumstances, including the novelty of the technology, risks and benefits associated with the use of the technology, scope of the representation, and sophistication of the client. Communication concerning the use of technology is required when the use presents a significant risk to the client's interests or materially affects the scope, cost, manner, or decision-making process of the representation. Routine, lawyer-directed use that does not present such a risk or material effect does not, standing alone, require separate disclosure under this rule.

Rule 3.3, Comment [3] - Verification and Separate Legal Requirements

[3] A lawyer's duty of candor towards the tribunal includes the obligation to verify the accuracy and existence of cited authorities, including taking reasonable steps to ensure no cited authority is fabricated, misstated, or taken out of context, before submission to a tribunal, including any cited authorities generated or assisted by artificial intelligence or other technological tools. Nothing in this comment independently requires disclosure of the use of technology to a tribunal unless disclosure is required by applicable law, rule of court, standing order, or tribunal directive. Any separate requirement concerning review, verification, certification, or disclosure must be satisfied according to its terms.

Rule 5.1, Comment [1] - Agentic Governance and External Requirements

[1] Paragraph (a) requires lawyers with managerial authority within a law firm to make reasonable efforts to establish internal policies and procedures designed, for example, to detect and resolve conflicts of interest, identify dates by which actions must be taken in pending matters, account for client funds and property, ensure that inexperienced lawyers are properly supervised, and govern the use of artificial intelligence, in accordance with the Rules of Professional Conduct. Policies governing the use of artificial intelligence should account for the technology's degree of autonomy, access to protected information and connected systems, ability to act or communicate externally, and capacity to affect substantive legal decisions. Reasonable measures may include authorization limits, access controls, source-traceability requirements, testing and monitoring, escalation procedures, and meaningful lawyer approval before material external communications, filings, advice, or other consequential actions. Reasonable firm policies should also identify and account for applicable statutory, court-imposed, client-imposed, and tribunal-specific requirements concerning technological assistance, including requirements governing protected information, verification, approval, certification, and disclosure.

APPENDIX B - AUTHORITIES AND MATERIALS CONSIDERED

A. State Bar and California Supreme Court Materials

1. State Bar of California, *Proposed Amended Rules of Professional Conduct 1.1, 1.4, 1.6, 3.3, 5.1, and 5.3* (clean and redline versions) (2026).
2. State Bar of California, *Proposed Amendments to the Rules of Professional Conduct Related to Artificial Intelligence*, initial and June 2026 modified public-comment materials (2026).
3. State Bar of California Standing Committee on Professional Responsibility and Conduct, *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* (updated May 14, 2026).
4. Letter from the California Supreme Court to the State Bar of California concerning artificial intelligence and agentic AI (Aug. 22, 2025).

B. Prior Public-Comment Materials

5. Jonathan Z. Morris, *Public Comment on Proposed Amendments Addressing Artificial Intelligence: AI, Legal Judgment, and the Future of Competent Lawyering* (Mar. 20, 2026).
6. COPRAC correspondence concerning the initial public comment and invitation to provide additional comment on the June 2026 Modified Proposal (2026).

C. Legislative Materials Concerning Senate Bill 574

7. Senate Bill 574 (Umberg), as amended July 2, 2026.
8. California State Legislature, legislative history and status materials concerning Senate Bill 574 (2025-2026 Regular Session).
9. Assembly Committee on Privacy and Consumer Protection, *Analysis of Senate Bill 574* (July 1, 2026).

D. Specific California Practitioner Commentary

10. Christian Puzder, *COPRAC's AI Warning Highlights Legal Industry's Accountability Gap*, Daily Journal (Apr. 2, 2026).
11. Eydith Kaufman, *Attorney-Client Confidentiality Melts Away When AI Enters the Room*, Daily Journal, Resolution Issue Supplement (Apr. 15, 2026).
12. Susie Dent, *New Guidance, Proposed Rule Changes on Generative AI in Legal Practice*, Daily Journal (June 4, 2026).
13. Daniel Butler, *California State Bar's Proposed AI Ethics Rules Put Attorneys on Notice*, Daily Journal (June 12, 2026).
14. Joe Stephens, *California's Proposed AI Ethics Rules Have a Blind Spot for Agentic AI*, Daily Journal (July 2, 2026).

E. Additional Practitioner and Judicial Materials

15. Additional recent California practitioner commentary concerning verification, agentic AI, accountability, auditability, confidentiality, privilege, client communication, and law-firm governance in connection with artificial intelligence (2026).
16. Recent state and federal judicial decisions, sanctions orders, standing orders, local rules, and other court materials concerning attorney, litigant, and judicial use of artificial intelligence, including matters involving fabricated or misstated authority, verification and candor, confidentiality and privilege, work-product protection, disclosure of AI use, and professional responsibility (2023-2026).