

Insurance Carriers, Panel Counsel and AI Controls

What carriers now require, what the insured's consent allows,
and how defense firms can show they follow the rules.



AN ICME LABS GUIDE · SEPTEMBER 2026 · ICME.IO

In brief

Carriers are beginning to add AI requirements to panel counsel guidelines. Some of them, especially records and outside-audit provisions, can create a confidentiality problem when meeting them would expose information about the insured's representation. This guide separates what the guidelines require today from a proposed evidence model that reduces that exposure.

- 1

Carrier AI rules are new and still rare.

Fulkerson Advisors found AI provisions in 20 of 1,054 published outside counsel guidelines. Two of the 15 insurers in the sample have them: The Hartford (2026) and Arch Insurance (2025).
- 2

The two published models differ.

The Hartford requires disclosure, records that allow AI use to be recreated, and audits by the carrier or its third-party auditor. Arch requires written approval before each AI use, senior review, records on request and an audit right.
- 3

Outside audits can require the insured's consent.

ABA Formal Opinion 01-421, the Montana Supreme Court and ethics opinions in at least seven states require the insured's informed consent before confidential information goes to a carrier's third-party auditor. AI prompts may contain far more of that information than a billing entry.
- 4

Courts now set their own AI rules.

Federal judges require certificates that AI output was verified and that no privileged information was disclosed. In 2026, several courts wrote AI restrictions into protective orders.
- 5

Evidence that reveals no content reduces the conflict.

A cryptographic receipt, issued when a firm's AI control runs, lets a carrier confirm that its guideline was checked on each AI action routed through the control, without an audit and without needing to obtain the insured's consent, because no prompt, output or file leaves the firm. It supplements a carrier's records and audit terms until the carrier agrees it satisfies them.

CONTENTS

03	Carriers are starting to write AI rules	07	What courts require, and what a firm can hand a carrier
04	Who pays, and the pressure on defense costs	08	A defense matter, stage by stage
05	The Hartford and Arch, side by side	09	One matter under a pre-approval guideline
06	Audit rights and the insured's consent	10	What carriers and defense firms can do now

Insurers have the highest observed share of AI clauses in the published sample, but the sample is small.

AI provisions in outside counsel guidelines first appeared in 2024. Fulkerson Advisors found them in 20 of 1,054 public guidelines, and in 2 of the 15 insurer guidelines it reviewed. The figures describe the sample, which is dominated by public bodies and universities, and do not estimate the carrier market.

EXHIBIT 1
Share of sampled public guidelines with an AI clause, by guideline date (2026 through September 16)





EXHIBIT 2
Insurers compared with all clients



EXHIBIT 3
What the 20 AI clauses require (number of guidelines)




Source: Fulkerson Advisors, What Clients Now Tell Their Law Firms About AI, September 2026. The sample is weighted toward public bodies and universities, and most carrier guidelines are not public, so the study cannot estimate the true carrier adoption rate.

Carriers have not settled who pays for defense AI, and cost pressure keeps rising.

Surveyed carrier executives expect AI to lower defense costs. Most have not been asked to fund it, and many have not been asked for permission to use it. If AI-assisted hours grow, carriers will have more AI time entries to check, and an audit cannot reach every matter.

EXHIBIT 4

Who should pay for law firms' AI tools, in a survey of senior claims and litigation executives



Sources: Insurance Journal, AI for the Defense: Should Insurers or Law Firms Pay? (April 20, 2026), reporting the 2026 CLM Litigation Management Study; sources listed under each figure.

The two published models share an audit right and differ on approval.

The Hartford conditions AI use. Arch requires written permission before each use.

EXHIBIT 5
Carrier AI requirements for defense firms

REQUIREMENT	THE HARTFORD (REVISED APRIL 20, 2026)	ARCH INSURANCE (2025)
Approval before use	Not required	Written consent "in each instance," naming the claim, attorney, tool and task
Disclosure	Name the AI program in each AI time entry; "Generative AI," "Gen AI" or "GAI" recommended	Through the approval request
Records	"Records (e.g., inputs, prompts, data and decisions) that allow for the recreation of generative AI usage"	"Detailed records" of approved AI use, available on request
Audit	By The Hartford "or its third-party auditor"	Arch "retains the right to audit"
Review	Verify the accuracy of all AI output and court submissions	Review by "a senior employee or shareholder"
Public tools	No nonpublic information in "any open-source AI system (e.g., ChatGPT)"	Covered by the consent requirement
Billing	Actual time only; no charge for AI licenses	No billing for both AI and human work on the same product
Breach	Defend and indemnify The Hartford for an actual or alleged breach	Self-report within two business days; indemnify Arch for any AI use

Under The Hartford's terms, a firm tags entries, keeps recreatable records and stands ready for an audit. Under Arch's, it requests approval matter by matter and cannot let any AI use run ahead of its approval. In both models, AI disclosure reaches the carrier as free text in a time entry. **We found no AI-specific field or UTBMS code in the LEDES and UTBMS materials we reviewed as of September 2026.**

Sources: The Hartford, Use of Artificial Intelligence, Claims Panel Counsel Case Management Guidelines (revised April 20, 2026); Arch Insurance, Litigation Counsel Guidelines (2025). Both are published for panel counsel and were verified October 3, 2026.


Outside audits can require the insured’s consent when they expose confidential information.

A defense firm’s confidentiality duty runs to the insured. Sharing with the carrier itself is often permitted, and some states treat the carrier as a co-client. Disclosure to a carrier’s outside auditor is a distinct question, and AI prompts may contain far more case information than a billing entry. These authorities predate generative AI; applying them to prompts is an analogy, not a holding.

EXHIBIT 6
The authority on carrier audits and the insured’s confidentiality

AUTHORITY	STATUS	HOLDING
ABA Formal Op. 01-421 (2001)	ABA advisory opinion	A lawyer “may not... disclose the insured’s confidential information to a third-party auditor hired by the insurer without the informed consent of the insured.”
ABA Formal Op. 08-450 (2008)	ABA advisory opinion	Counsel must withhold from the carrier information relating to the representation that is damaging to the insured’s interests under the policy.
In re Rules of Prof’l Conduct, 2000 MT 110	Montana Supreme Court decision	Disclosure to third-party auditors without the insured’s “contemporaneous fully informed consent” violates confidentiality. Consent given in advance through the policy is not enough.
Alabama, Mississippi, North Carolina, Oklahoma, Washington, West Virginia and Texas ethics opinions (1998 to 2000)	State ethics opinions	Each requires the insured’s consent before bills go to a carrier’s auditor. Texas 533 also bars agreeing to restrictions that interfere with independent judgment.
ABA Formal Op. 512 (2024)	ABA advisory opinion	Applies existing duties to generative AI. Whether client information may go into a tool, and whether consent is needed, depends on the tool and the circumstances. Lawyers must disclose AI use when a client’s outside counsel guidelines require it.

EXHIBIT 7
Two ways AI evidence can reach the carrier



Opinion 512 makes a client’s guidelines one trigger for disclosure.

It imposes no blanket duty to disclose AI use. The insured, as the represented client, holds the consent that governs what an outside auditor can see.

The Hartford’s terms show the conflict.

Its records clause covers “inputs, prompts, data and decisions,” and its audit clause extends to a third-party auditor. Its billing guidelines also say it does “not expect firms to follow any billing or litigation management guidelines that conflict with ... ethical directives.” A receipt that discloses no content avoids the disclosure problem; it satisfies the records and audit terms only if the carrier agrees.

Sources: ABA Formal Opinions 01-421, 08-450 and 512; 2000 MT 110; the state opinions named above; The Hartford, Billing Guidelines and Use of Artificial Intelligence (2026).

Courts and carriers increasingly overlap on verified work and protected information.

Some judges now require verification of AI-assisted filings, disclosure of AI use, or certification that AI use exposed no confidential material. Damien Charlotin’s independent database had identified 2,077 decisions on AI-hallucinated filings as of September 23, 2026, 825 involving lawyers.

EXHIBIT 8
Court requirements that bear on AI use in defense work

AUTHORITY	WHAT IT REQUIRES
Judge Marcia A. Crone, E.D. Tex., Standing Order MC-11 (July 2026)	Certify that counsel “independently verified the accuracy” of AI-assisted content and that AI use “has not resulted in the disclosure of confidential or privileged information.”
Morgan v. V2X (D. Colo. Mar. 30, 2026)	In that case’s protective order, no confidential information may go into an AI tool unless the provider is barred from training on it or disclosing it, with deletion terms.
Jeffries v. Harcros (D. Kan. Mar. 25, 2026)	In that action’s amended protective order, no discovery materials may go into open AI tools.
United States v. Heppner (S.D.N.Y. Feb. 17, 2026)	Documents the defendant created on his own with a public AI service, not at counsel’s direction, were held neither privileged nor work product on those facts.

EXHIBIT 9
What a firm can hand a carrier, and what it exposes

		CAN THE CARRIER CHECK IT?	SATISFIES THE GUIDELINE AS WRITTEN?	WHAT IT EXPOSES
1	Firm AI policy States what the firm intends to do.	No	Partly; The Hartford requires one	Nothing
2	Tagged time entry The firm’s own note of the tool and task.	On the firm’s word	Yes; both require disclosure	Tool and task
3	Prompt and input records The firm’s log of what AI was given.	By reading them	Yes; both require records	Facts of the case
4	Audit A review after the work is done.	After the fact	Yes; both reserve the right	Prompts, outputs, files
5	Cryptographic receipt Issued when the control runs, before the action.	Issuer and integrity, independently	Only if the carrier agrees	Rule, result and time

Sources: orders and decisions named above; Damien Charlotin, AI Hallucination Cases Database (September 23, 2026). A written AI policy did not prevent a \$59,500 sanction against a national insurance defense firm and its partner in Cook County in December 2025 (Chicago Sun-Times); a policy alone does not establish compliant execution. A receipt lets the carrier verify who issued it and that it was not altered; the facts it records still depend on the control’s design.

Every stage of a matter has questions a control can decide and questions only the lawyer can judge.

Carriers write the yes-or-no rules into guidelines, and courts write them into orders. A control can enforce those rules before an AI action runs and issue a receipt each time. It governs only actions routed through it, so firms must close off AI use outside it. The judgment is why the carrier hired the firm.

EXHIBIT 10
A defense matter, stage by stage

STAGE	WHAT A PRE-ACTION CONTROL DECIDES AND PROVES	WHAT THE LAWYER JUDGES
Assignment	The carrier's AI rules for this claim are loaded before any work starts	Whether AI suits the matter at all
AI approval	No AI use runs without an approval that matches the claim, tool and task	What to ask the carrier to approve
Research	Only approved tools run; no insured information goes to a public tool	Whether the authority is right and applies
Drafting and filing	Nothing is filed until a named attorney approves that exact version	Whether the argument is sound and the draft is ready
Discovery	Produced documents go only to tools the protective order allows	Relevance, privilege calls and strategy
Reporting and billing	Each AI time entry carries a receipt for the checks that ran	What to tell the carrier about the case

How the controls work

- 1

Lawyers write the rule in plain English.

"No AI use on a carrier matter without written approval for this claim, tool and task."
- 2

The rule is translated into formal logic, and a lawyer approves the translation.

The approved logic is versioned so a receipt can point to it. A solver checks each AI action routed through the control before it runs; the same inputs return the same answer every time. ICME Labs' Preflight is one implementation.
- 3

Each check issues a cryptographic receipt.

A tamper-evident record of the rule, the result and the time, bound to a fingerprint of the action. The carrier verifies that the firm's control issued it and that it has not been altered, without the prompt, the output or the file. It is evidence of the check, not proof that the facts behind it were true.

Disclosure: ICME Labs publishes this guide and builds Preflight. Statements about product capabilities are ICME's own unless independently sourced.

One matter under a pre-approval guideline.

A simulation of a matter under a guideline like Arch’s. Names, claim number and receipts are illustrative.

Rule: No generative AI may be used on a carrier matter unless the carrier has approved that use in writing for this claim, this tool and this task.

1 · REQUEST

Approval recorded

The carrier approves an enterprise tool for drafting discovery responses on claim 24-0917 and issues the approval ID.

2 · BLOCKED

Unapproved task stopped

An associate asks the tool to summarize medical records. The control blocks it before it runs.

3 · ALLOWED

Approved task runs

The claim, tool and task match the approval. The draft proceeds and a receipt issues.

PREFLIGHT RECEIPT · R-7F3A	
Rule	carrier AI approval
Approval ID	AP-24-0917-03
Result	Allowed
Checked	before the action ran
Content disclosed	none

THE INVOICE LINE THE CARRIER SEES

2.4 hrs · Draft responses to plaintiff’s first interrogatories · GAI: approved enterprise tool · Receipt R-7F3A

✓ Verified by the claims professional. No prompt, output or document received.

The approval ID comes from the carrier, so the fact the check relies on is not one the firm can create. For that to hold, the approval should be carrier-issued and bound to the claim, tool and task, with an expiry, so it cannot be reused for a different action. Under a disclosure-and-audit guideline like The Hartford’s, the check runs on where the data goes: it confirms the tool is an approved enterprise system before any nonpublic information reaches it.

14

AI actions checked

2

stopped before they ran

12

cleared

9 of 9

AI time entries with a receipt

WHAT A RECEIPT DOES NOT SHOW

✗ That the facts the check relied on were true

✗ That the lawyer read the document; it shows a named attorney approved that exact version

✗ What an AI vendor does with data after it leaves the firm

✗ Whether the legal work was good

✗ That every AI use passed through the control; use outside it leaves no receipt

✗ That the logic matches the guideline’s text; a lawyer approves that translation

QUESTIONS THE MARKET HAS NOT ANSWERED

? We found no billing code that identifies AI-assisted work

? We found no public example of a carrier accepting receipts in place of audit access

? We found no court ruling on consent for audits of AI records

? Malpractice carriers ask about AI; we found no premium credits and no AI exclusion on a lawyers’ policy form

What carriers and defense firms can do now.

CLAIMS AND LITIGATION MANAGERS

Verification on every matter instead of an occasional sample, at lower cost than an audit, without taking on the insured’s confidential information or needing to obtain the insured’s consent.

- 1 Ask which AI tools the firm uses on your matters and who approved each one.
- 2 Ask how the firm stops AI use your guidelines don’t allow before it happens.
- 3 Ask for evidence on AI time entries that you can verify yourself.
- 4 Ask what an audit of AI use would disclose about the insured, and how consent would be obtained.
- 5 Add the model clause below to your panel guidelines.

DEFENSE FIRMS

One control framework can run carrier-specific and matter-specific rule sets, keep prompts and files inside the firm, and support a lawyer’s court certification with system evidence.

- 1 List each carrier’s AI terms: approval, tagging, records, audit and public-tool limits.
- 2 Map which AI actions each term covers: research, drafting, summarizing records, handling productions.
- 3 Set a consent process before an auditor asks: what an AI audit would disclose, and how the insured decides.
- 4 Route AI actions through a check loaded with each carrier’s rules, so an unapproved use stops before it runs, and close off use outside it.
- 5 Reconcile before billing: every AI time entry carries evidence the carrier can verify.

A PROPOSED ALTERNATIVE EVIDENCE CLAUSE, WHICH REQUIRES CARRIER ADOPTION

Evidence of AI controls. Counsel shall operate controls that check each use of generative AI on Carrier matters against these Guidelines before the use occurs. For each AI-assisted time entry, Counsel shall provide a verifiable record of the rule checked, the result and the time, without disclosing prompts, outputs or other information relating to the representation of the insured. Carrier, and any auditor acting for Carrier, will review AI use through these records, and will request prompts, inputs or outputs only with the insured’s informed consent obtained at the time of the request.

The clause is directionally consistent with the governance and third-party oversight principles of the NAIC Model Bulletin on insurers’ use of AI, adopted in 26 jurisdictions as of August 31, 2026. Whether a defense firm’s drafting or discovery AI falls within the bulletin’s scope is a separate question.

See a receipt verified against a sample carrier guideline. Write to houman@icme.io.

About ICME Labs. ICME Labs builds Preflight, which checks each AI agent action against an organization’s rules before it runs and issues a cryptographic receipt a third party can verify without seeing the underlying data. ICME Labs is Stanford CodeX-affiliated. icme.io