

Fiduciary-Grade AI for in-house legal

A new standard for trust, verification, and defensibility

Thomson Reuters™

It's no surprise that many lawyers have experimented with consumer AI tools to save time on that day-to-day work.

The pressure on today's general counsel isn't just to support the business; it's to help the business make the right moves confidently. That means creating capacity to deliver timely strategic advice without compromising the day-to-day work the business depends on.

It's no surprise that many lawyers have experimented with consumer AI tools to save time on that day-to-day work. But when speed comes at the expense of trackable citations, confidentiality, or defensible reasoning, the shortcut can create more risk than it removes.

Legal teams need AI built for the standards of legal practice, not the standards of the general public. That's where Fiduciary-Grade AI™ comes in. This is AI built not simply to generate answers, but to support the professional responsibilities of legal practice.

The risk of consumer AI in high-stakes legal work

Consumer AI may be appropriate for brainstorming or drafting a vacation itinerary.

Legal professionals, however, operate in a very different environment. They're accountable for the advice they give, managing the company's risk, protecting confidential information, and are increasingly expected to demonstrate that AI-assisted work is used responsibly.

That's not just a matter of professional judgment. Bar associations have begun issuing formal AI ethics guidance. Governments are adopting AI governance requirements. Companies are establishing policies that restrict the use of consumer AI tools with sensitive information. Together, these developments are raising the bar for how legal professionals select, govern, and use AI.

As expectations continue to rise, the question is now whether the AI tools lawyers choose meet the standards their profession demands.

Fast and (often) free consumer-grade AI tools have several drawbacks.

First, consumer AI systems are designed to generate plausible responses, not authoritative legal analysis. While impressive for many everyday tasks, they can hallucinate facts, misstate legal authority, or rely on outdated or non-authoritative sources.

Their reasoning operates as a "black box." Users don't know where the answers come from, and they can't effectively validate the results. This creates liability and audit risks.

Some consumer AI tools may use submitted data for model training. This could inadvertently waive privilege over legal advice given to the business, or expose confidential corporate, employee, or customer information.

Most publicized AI failures have involved litigators because court filings become public. This does not mean transactional attorneys should be complacent. Misuse of consumer AI in-house is more likely to remain inside the organization and potentially go undiscovered until a crisis arises from a contract dispute, policy change, or government action. That means risk will accumulate over time and could expose the business to significant, unpredictable legal and business outcomes.

Defining Fiduciary-Grade AI

Fiduciary Grade AI is the standard from Thomson Reuters® for how AI should work in high stakes professions, including legal. The standard has four principles: to be Fiduciary-Grade AI, a tool must draw on authoritative, domain specific content; be built around rigorous privacy and security safeguards; be shaped and tested by human subject matter experts; and be designed to produce transparent outputs that can be verified.

The four pillars of Fiduciary-Grade AI

Fiduciary Grade AI is defined not just by what it produces, but by what it is allowed to access, retain and rely upon in generating outputs that inform professional judgment. There are four key pillars to assess an AI tool as Fiduciary-Grade or not:

AI grounded in authority; with access to the right context

A Fiduciary-Grade AI system must derive its substantive outputs from authoritative, curated, and domain-specific content, not just information scraped from the open internet. It also must operate with the full context required to complete professional work. Every material output must be traceable to a source that a qualified professional can independently locate, cite, verify, and trust. And only when AI agents can access, know, and act on the specific data, knowledge, systems, and tools can they complete the complex, multi-step tasks that professional work demands.

For example, an in-house lawyer assessing a global product launch needs AI grounded in current, authoritative law and informed by the company's own policies, contracts, and risk tolerances. A generic answer about data privacy is not enough. The lawyer needs analysis tailored to the product, jurisdictions, and business context, with sources that can be reviewed.

Data privacy and security are imperative

Where privacy is paramount, Fiduciary-Grade AI is built to protect it. Privacy and security must be structural features of the system's architecture, not policy overlays or configurable options.

For example, an in-house lawyer using AI to organize an internal investigation may be working with employee records, interview notes, allegations of misconduct, and privileged legal analysis. Fiduciary-Grade AI must protect that information by design, with clear controls over access, retention, and use. It's not enough to rely on the user to avoid entering sensitive details.

The four pillars of Fiduciary-Grade AI aren't theoretical goals or potential future capabilities; they're central to how CoCounsel Legal was built and is being updated. Learn more now.



Built with human expertise, not just human oversight

Professional workflows must be designed, tested, and continuously refined with meaningful involvement from credentialed subject matter experts in the relevant professional domain. When ambiguity or risk arises, the system must recognize its limits and bring professionals back in rather than generating an output that overstates its reliability. This keeps accountability human and outcomes defensible. Fiduciary-Grade AI requires that users have access to real-time human support to ensure transparency and trust.

For example, an AI tool reviewing a customer contract may identify that a limitation-of-liability clause departs from the company's approved playbook. Rather than confidently recommending acceptance or rejection, it should flag the deviation, explain the legal and commercial considerations, and route the issue to counsel for human judgment.



Transparent, verifiable reasoning

By clearly surfacing and referencing the sources it relies on, AI must be able to provide a reviewable trail of what the system did and what it relied on, sufficient to allow a qualified professional — and, where applicable, a regulator, court, or auditor — to evaluate the basis for the output and determine whether the result is reliable and defensible. Making each step in its planning, reasoning, and execution process visible to the user is vital to verifying the AI has reached the correct answer.

For example, when advising executives on the impact of a new regulation, an in-house lawyer needs to see the authorities behind the analysis, the assumptions applied, and the steps connecting the law to the recommended business response. That reviewable trail allows counsel to test the work, explain it to leadership, and defend the company's approach if it is later questioned by a regulator or auditor.

The four pillars of Fiduciary-Grade AI aren't theoretical goals or potential future capabilities; they're central to how CoCounsel Legal was built and is being updated. Learn more now.

Maintaining human accountability

Fiduciary-Grade AI is designed to support professional judgment, not replace it. Lawyers remain responsible for their work. AI can accelerate research, summarize information, or draft work product, but it cannot assume the professional responsibilities that accompany legal practice.

That means Fiduciary-Grade AI should recognize the limits of automation. When a question depends on missing facts, novel legal issues, or significant business risk, the system should surface uncertainty, explain why professional judgment is required. Then it can bring the lawyer back into the decision-making process.

For an in-house legal team, that might mean

- Identifying that a customer contract contains non-standard indemnification language requiring legal review
- Recognizing that a new regulation has unresolved interpretive questions
- Flagging that an employment issue depends on facts not contained in the prompt

In each case, the AI helps lawyers focus their expertise where it adds the most value instead of encouraging them to rely on automation alone.

Maintaining human accountability also extends beyond the software itself. Fiduciary-Grade AI should be developed and continuously refined with meaningful involvement from experienced legal professionals, and users should have access to knowledgeable human support when questions arise about the system's outputs, underlying sources, or appropriate use.

After all, the objective isn't simply to make legal work faster. It's to help lawyers spend more of their time doing the work only they can do: exercising judgment, guiding business decisions, and earning the trust that makes them valued strategic advisors.

Navigating global regulatory compliance

As organizations adopt AI, regulators and professional bodies are making it clear that accountability remains human. Emerging AI governance frameworks consistently emphasize many of the same principles. These include human oversight, transparency, data protection, security, and the ability to explain and defend AI-assisted decisions.

For legal professionals, these expectations reinforce obligations that have long existed. The American Bar Association has reminded lawyers that using generative AI does not diminish their duties of competence, confidentiality, supervision, or independent professional judgment. In the United Kingdom, the Solicitors Regulation Authority has similarly emphasized that solicitors remain responsible for work produced with AI tools. Meanwhile, the European Union's AI Act establishes new governance expectations around transparency, risk management, human oversight, and documentation for many AI systems used in professional contexts.

Fiduciary-Grade AI isn't an abstract framework. It should shape how legal AI is designed, the information it relies on, and the way it supports lawyers throughout their work.

Fiduciary-Grade AI is designed with these principles in mind. Rather than treating governance as an afterthought, it incorporates authoritative legal content, strong privacy protections, human accountability, and transparent, reviewable outputs into the foundation of the system. Those capabilities align with the governance principles increasingly reflected in legal ethics guidance and AI regulation.

Fiduciary-Grade AI supports regulatory compliance as technology evolves

Regulatory trend	Fiduciary-Grade AI response
Human oversight	Lawyers remain accountable; AI recognizes its limits and escalates when appropriate.
Transparency	Outputs are grounded in authoritative sources that can be reviewed and verified.
Privacy and confidentiality	Sensitive legal information is protected through enterprise-grade safeguards.
Governance and documentation	AI is designed to support explainable, defensible professional workflows.

CoCounsel Legal: Fiduciary-Grade AI in Practice

Fiduciary-Grade AI isn't an abstract framework. It should shape how legal AI is designed, the information it relies on, and the way it supports lawyers throughout their work. CoCounsel Legal from Thomson Reuters was built around these principles. It helps in-house teams accelerate routine work while preserving the professional judgment and accountability at the heart of legal practice. By combining authoritative Thomson Reuters content with agentic AI that can plan and execute multi-step workflows, and an organization's own documents, CoCounsel helps legal departments deliver faster, more consistent guidance without compromising professional standards.

Grounded in authoritative legal content

CoCounsel's research capabilities are grounded in trusted legal content, including Westlaw and Practical Law. It is designed to surface the authorities behind its analysis. Whether an in-house lawyer is evaluating the impact of a new regulation or preparing guidance for a product launch, the goal is to provide answers that can be reviewed, verified, and defended.

Legal teams use CoCounsel to provide executives with timely market-entry guidance, accelerate procurement reviews, and support compliance across multiple jurisdictions, all without increasing headcount.



Built to support legal workflows

CoCounsel's guided workflows plan and execute multi-step tasks such as research, document review, and drafting while bringing lawyers back into the process when decisions require professional judgment. Prompt libraries developed by legal experts help teams standardize common work while reducing the need to start from scratch.



Designed to create capacity for strategic work

By accelerating contract review, due diligence, policy drafting, and legal research, CoCounsel helps legal departments spend less time on repetitive work and more time advising leadership, enabling business initiatives, and managing risk. Legal teams use CoCounsel to provide executives with timely market-entry guidance, accelerate procurement reviews, and support compliance across multiple jurisdictions, all without increasing headcount.

The legal team at Century Communities adopted CoCounsel and found that it created significant efficiencies. "There's way more to it now, but the genesis for me was the ability to take a ton of data and summarize it quickly and then spit out a link to where it found that data so you can do your own double checking," says Century Communities General Counsel Jarret Coleman. "A task that would previously have taken an hour was completed in five minutes or less."

CoCounsel doesn't replace the trusted advisor. It helps create more opportunities for lawyers to act as one.



Built for trust

Perhaps most importantly, CoCounsel reflects the principle that lawyers remain accountable for their work. Its outputs are transparent and reviewable, allowing lawyers to verify the underlying authorities before relying on them. That design philosophy mirrors the core principles of Fiduciary-Grade AI: authoritative content, human accountability, and transparency working together to support professional judgment rather than replace it.

Reaching a high bar for legal work

"Lawyers don't want to just operate software, and that's not what great AI should do," says Andrew P. Medeiros, managing director of Innovation at Troutman Pepper Locke. "CoCounsel keeps them in the analytical mindset they were trained for: going back and forth, challenging answers, and steering the work."

Fiduciary-Grade AI doesn't lower the standard for legal work. It helps in-house counsel meet that standard more efficiently, creating capacity to do what the business, executives, and boards increasingly expect of them: provide trusted judgment that helps the business move with confidence.

Medeiros, Coleman, and other legal professionals on CoCounsel are benefiting from Fiduciary-Grade AI. For in-house legal departments, only Fiduciary-Grade AI provides the authoritative grounding, privacy, expert-backed design, and transparent reasoning required to deploy AI safely and effectively.

Discover how CoCounsel Legal delivers Fiduciary-Grade AI for your in-house legal team and request **a demo of CoCounsel Legal**.

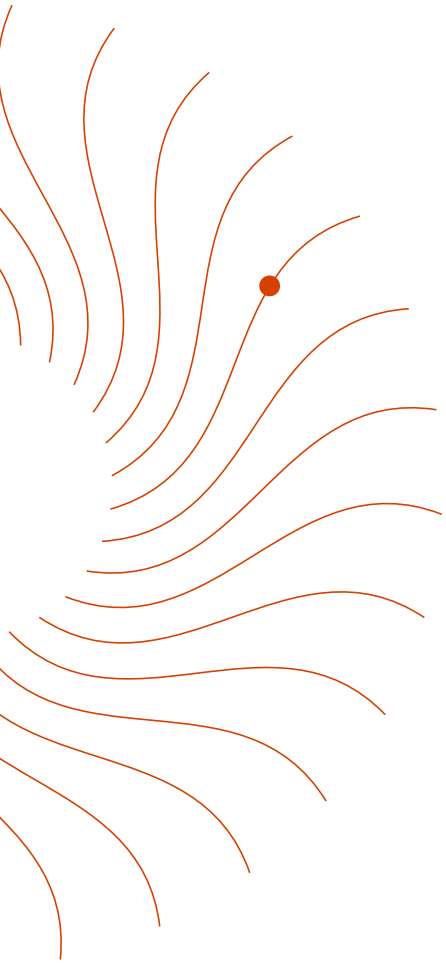

© 2026 Thomson Reuters. Fiduciary-Grade AI is a trademark of Thomson Reuters. All rights reserved. TR8479938/9-26

Thomson Reuters™