Why legal AI adoption slows after pilots

Written by Damien Riehl

As featured in Artificial Lawyer

If we've experienced nothing else over the past three years, we've certainly heard a familiar question that continues to surface regularly: *what's our AI strategy?* The phrasing implies a bounded decision on tools. But in practice, it's mostly organizational.

Public discussion around legal AI still often focuses on model performance, hallucination rates, or drafting quality. Inside firms, progress tends to depend on something more operational. A system can produce a credible answer, but how might the firm integrate that capability to its legal work — for that practice group, by that lawyer.

That distinction usually shapes firm adoption. How an organization's workflows, controls, and responsibilities evolve often matters more than the specific software chosen. That becomes obvious the moment a pilot converts to work on real, live matters.

Early pilots might establish confidence. They can show that the system can handle information appropriately, and its output is usable (and sometimes even stunning). Once the same capability sits inside an active matter, attention might shift to how the firm works with it: where matter context appears in the workflow, who stands behind the result, and how its use fits alongside systems, staffing, billing, and client expectations. Context is everything: true for people, true for workflows, and true for AI.

Where early momentum meets operational reality

Implementation tends to follow predictable patterns. Firms that begin by clarifying how information is governed and logged can expand more deliberately. Only after establishing those foundations might attention turn to daily workflows. Adoption grows when lawyers can use familiar systems to run matters. When adoption requires stepping outside familiar systems, usage could remain intermittent, regardless of technical quality. Of course, what is "familiar" is a moving goalpost. Every "familiar" tool was formerly "strange." Usage accelerates familiarity.

The legal market moves similarly: Hemingway wrote that bankruptcy happens "[g]radually and then suddenly." With legal AI, we might be entering Phase Two. AI has migrated beyond "niche destination" — it's increasingly appearing within research tools, document environments, and matter management platforms. Capabilities such as Vincent within

broader practice management environments (e.g., Business of Law, Practice of Law) illustrate this shift toward embedding.Functionality moves from “additional steps” friction to becoming “part of my workflows.”

After initial evaluation, familiarity leads to repeated use. Proficiency can grow because outputs behave consistently. And in professional settings, deterministic consistency (aka reliability) reigns supreme.

The role of operational clarity

As firm usage expands, operational questions can surface. Firms need a shared understanding of data handling, retention, and activity records before treating the output as its work product. These considerations are substantive, not technical, yet they determine whether the capability spreads beyond small groups.

The same dynamic appears at scale. A tool that feels straightforward for one lawyer might introduce questions across practice groups. So alignment around the collective process might become as important as the underlying capability. Where alignment exists, use can grow naturally. Without that alignment, adoption might remain cautious (or stalled) even when results are strong.

Once capabilities are embedded and routinely used, the problem might change again. Conversations might move from implementation, migrating to performance. Firms may less-frequently ask whether lawyers will use it — and more often assess whether it’s improving matters in ways that the partnership and clients can see.

When the capability enters routine matter workflows

Over time, decisive factors are often architectural. Data location, auditability, and workflow integration determine whether usage is flat or expands. When those elements align, changes inside firms feels incremental. Turnaround times shorten, reporting simplifies, and working practices adjust to business as usual: the new normal. .

But from the outside, firms might appear to “suddenly” adopt AI. An overnight success… months or years in the making. What looks like a recent decision is usually the point where earlier groundwork finally accumulated enough traction and trust to support routine use.

So taking time to properly establish legal AI’s foundation shouldn’t be viewed as hesitation or delay. It’s closer to building infrastructure: largely invisible while it happens, and only obvious once complete. And self-evident when it’s missing. In that infrastructure period, firms aren’t really exploring whether AI works; they’re more often deciding if and how lawyers and legal professionals can stand behind the tool’s outputs.

Once adopting a tool routinely, firms can then focus on outcomes. Where time is actually saved. Where friction and risk are reduced. Where quality improves in ways that withstand client scrutiny. That's when vibes turn to measurement — and when policy governance favors operational discipline.

When everyone stops discussing the technology, instead assessing results, only then do organizations feel adoption. And what follows is quieter and more consequential: the well-earned benefits of making the capability dependable, measurable, and worth keeping.

About Clio

Clio is the global leader in legal AI, powering legal work across the profession, from law firms of every size to corporate legal teams, government agencies, and courts. Its Intelligent Legal Work Platform unites the business and practice of law, connecting the people, information, and processes behind legal work. Clio's AI brings the context together with legal sources drawn from Clio Library, a growing collection of more than one billion legal documents. Over 400,000 legal professionals in more than 130 countries rely on Clio to run their legal world. Learn more at www.clio.com.