

We Gave a 12-Year-Old a Shotgun and Said “Figure It Out”: Inside Legalweek’s Most Honest Panel on Agentic AI

By Matthew A. Mishak, Esq. | Founder & CEO, LegalTek.ai | March 11, 2026

Filed from Legalweek New York 2026 — Javits Center, New York City

That’s how one panelist described where the legal profession stands with AI right now.

Not the polished, aspirational version you’ll hear from most vendor booths on the exhibit floor. The real version. The one where four of the sharpest minds in legal AI sat down at the “AI Agents 101: What They Are and Why They Matter” panel and, for sixty minutes, told a packed room of attorneys, legal ops professionals, and technologists exactly how unprepared we all are for what’s coming—and why that should terrify and excite us in equal measure.

The panel, moderated by **Julie Chapman**, Head of Legal for North America at LexisNexis, featured **Ethan Wong**, Product Counsel at Anthropic; **Kassi Burns**, Counsel at King & Spalding; and **John Rosenthal**, Partner at Shook Hardy & Bacon. The session was supported by LexisNexis as part of Legalweek’s AI Agents and Automation track.

Here’s what happened—and why every lawyer building, buying, or even thinking about AI should be paying attention.

Forget “AI Agent.” Think Workflows.

The panel opened with a deceptively simple question: *What is agentic AI?*

Ethan Wong of Anthropic broke it into two categories. The first is workflow-based agents—systems that follow prescribed rules. If X happens, then do Y. The second is autonomous agents—systems that exercise independent judgment when given a general task without specific rules. Most of what the legal industry is calling “agentic AI” right now, Wong explained, falls squarely in the first bucket.

Kassi Burns of King & Spalding brought a musician’s sensibility to the definition, framing

agentic AI as *orchestration*. It's not one model doing everything, she said. It's multiple foundation models, each handling a sub-component of a larger workflow, with an orchestration layer managing the whole process. "Thinking really big picture," Burns explained, "agentic AI systems aren't going to be one single thing that does everything for you. It's identifying sub-components and optimizing them."

Then John Rosenthal went off-script in the best possible way.

"There's a lot of market confusion created by a lot of marketing," he said, cutting through the noise that permeated the exhibit hall one floor below. He walked the audience through the actual state of the technology: we're at the very front end, still stuck at the prompt window, with agents that are "rather simplistic." But the trajectory? "In the next two to three years," Rosenthal predicted, the industry will move from single-prompt systems to multi-agent workflows that are pre-built and ready to deploy.

His morning workflow vision was compelling: press a button, say "good morning, launch my BD," and have ten different agents pull from news sources, case law databases, and Wall Street Journal articles to produce a personalized business development briefing with actionable next steps.

Burns took it even further: "You could have it monitor you as you sleep, and maybe it will make some assessment... you're going to adjust what gets delivered." The room laughed, but nobody dismissed it.

The Weekend Warriors Are Building the Future

Perhaps the most actionable part of the session was the panelists' candid discussion of how they're personally experimenting with these tools—not in their professional capacity, but on their own time.

Rosenthal, a senior partner at an Am Law firm, spends his weekends building agentic workflows. Two weeks before the panel, he built a case law update agent in ninety minutes that scans legal resources and produces daily summaries. The week before Legalweek, he created an agent that scans his mailbox and desktop to auto-generate timesheets, complete with import links to the firm's billing system.

Burns shared her experience with Zapier, building automated workflows that monitor specific YouTube creators and send email summaries when certain topics are discussed. She emphasized that the exercise isn't just about the output—it rewires how you think. "It forces your brain to work a little bit differently, because it does force you to break things off into components."

Wong shared the example of a friend in finance who coded a restaurant reservation agent over a single weekend using Claude Code—a system that monitors for openings at hard-to-book restaurants and pings him when a slot appears. “It didn’t work perfectly,” Wong admitted, “but that’s kind of the space I would love to get to in my personal life.”

The implicit message was unmistakable: if you’re waiting for someone to hand you a turnkey agentic solution, you’re already behind.

The Compliance Time Bomb Nobody’s Talking About

Then the conversation turned darker—and, in my opinion, far more important.

Rosenthal channeled his inner *Lost in Space* robot: “Danger, Will Robinson.” His concern? The proliferation of agents across enterprises without any governance wrapper. We’re going to create these agents, he warned, and then nobody will know whether they’re up to date. Nobody will know who created the agent. Nobody will know, ten years from now, whether the agent is still compliant. Nobody will know whether there’s a privacy violation embedded in it, or whether it’s accessing data it shouldn’t.

And it’s not just a legal department problem. “Think of our clients,” Rosenthal pressed. “They don’t make briefs. They make automobiles. They brew beer. They build software. All of the business people are doing the same thing we’re doing—saying, ‘How do I get an agent to make my life simpler so I can make the product quicker and cheaper?’ They’re doing the same thing with the same danger.”

Rosenthal dropped a statistic that should keep every managing partner up at night: only about 50% of Am Law 100 firms have AI policies in place. Almost none have published AI guidelines for their clients.

Burns connected the dots to data hygiene—a foundational issue that nobody finds exciting but that will determine whether agentic systems succeed or become enterprise liabilities. “Garbage in, garbage out” isn’t just a cliché when you’re chaining together multi-step automated workflows that touch sensitive client data across multiple platforms.

Wong, speaking from the product side at Anthropic, added the design perspective. When Claude has access to a user’s entire desktop through Cowork, there are intentional friction points built into the experience. Before Claude deletes something, for example, it pauses and asks for confirmation. These aren’t bugs in the user experience—they’re deliberate guardrails. The question is whether the legal industry and its clients will build similar safeguards into their own implementations.

The Question I Had to Ask

Full disclosure: I was in the audience, and I couldn't stay silent.

My question was about the elephant in the room—the gap between where the technology companies want guardrails and where the law actually provides them. We've seen it play out in real-time with the Anthropic-Pentagon FASCSA dispute, which I've been writing about extensively. Do you let an AI agent make life-or-death decisions? And when the technology is racing ahead of any legislative framework, how do we navigate that void?

Rosenthal's answer was characteristically blunt: the providers want legislation because they want a safe harbor. They don't have one. Products liability cases are starting to flow—claims that AI provided harmful medical advice, legal advice, or exposed minors to dangerous content. Without legislative protection, the providers are exposed, and so is every organization deploying their tools.

His prescription? Go back to policy and training. But not the “paper tiger” version that most organizations currently have. Practical, operational guidance backed by actual compliance infrastructure. And critically, he argued, this work needs to move out of the legal department silo and into compliance.

The Hallucination Problem Isn't Going Away

Rosenthal raised what might be the most uncomfortable truth of the entire panel: generative AI can hallucinate up to 50% of the time in certain applications, and we're building agentic platforms on top of that foundation.

“We're stuck in workflows, and to some extent, we're going to become numb,” he said. He drew a parallel to Google's AI-generated search summaries—often inaccurate, yet increasingly treated as authoritative. The legal profession, he warned, will face the same problem at enterprise scale unless firms build validation and backstopping into every agentic workflow.

Burns agreed, pointing to the promise of an orchestration layer that includes human-in-the-loop escalation points and feedback mechanisms. The question isn't whether to keep humans involved—it's defining *how much* human oversight is sufficient.

The Access-to-Justice Paradox

The panel's final exchange touched on a tension that will define legal AI for the next decade:

the unauthorized practice of law versus access to justice.

Wong noted that Anthropic's usage policy explicitly prohibits Claude from providing legal advice—when asked legal questions, it directs users to consult an attorney. But Rosenthal pointed out that reality is outpacing policy. An OpenAI case just last week found that despite explicit disclaimers, an AI system was effectively providing legal advice.

Burns raised the access-to-justice angle: for people who can't afford an attorney, AI represents a potential lifeline. But how does that potential benefit coexist with unauthorized-practice-of-law concerns? "We'll have to wait and see," she said. But nobody on the panel seemed confident that the current frameworks are adequate.

What This Means for Legal Tech Builders

I left that panel more convinced than ever that what we're building at SilverTung is on the right track.

Here's what the panel confirmed for me:

Workflow-first architecture wins. Every panelist emphasized thinking in workflows, not tasks. When Wong described Claude's skills system—step one, do this research; step two, look at these cases; step three, pull these facts—he was describing exactly the kind of structured, sub-component approach that SilverTung was designed around from day one. Ohio domestic relations practitioners don't need a general-purpose AI. They need a system that understands the workflow of generating an Affidavit 1, populating a Shared Parenting Plan, or drafting a Separation Agreement—with guardrails built in at every step.

Governance can't be an afterthought. Rosenthal's compliance warnings aren't theoretical—they're operational. This is exactly why we built the COUNSEL framework and the G3M governance model into SilverTung's DNA. When your AI system is producing documents that get filed with a court, "move fast and break things" is malpractice.

The human-in-the-loop isn't going anywhere. Wong's description of Anthropic's design philosophy—pausing before destructive actions, asking for confirmation, keeping the user in control—mirrors what we've built with our Tattle-Tale compliance oversight system. The attorney always has the final say. Always.

Data quality is destiny. Burns's emphasis on data curation and hygiene validates our data-first approach with the PostgreSQL schema we've built. When you're pulling financial data through Plaid, mapping it to Ohio Affidavit 1 fields, and feeding it through an AI-driven analysis engine, the quality of that pipeline determines everything downstream.

The Bottom Line

We are at the very beginning of a revolution in legal practice. The panelists at Legalweek 2026 were refreshingly honest about that—about the promise, the peril, and the vast gulf between where we are and where we're going.

The lawyers who will thrive aren't the ones waiting for a perfect, turnkey solution to appear on the exhibit floor. They're the ones spending their weekends—like John Rosenthal—building workflows, testing tools, and figuring out where the failure points are before their clients discover them the hard way.

The firms that will lead aren't the ones with the flashiest AI announcements. They're the ones building governance frameworks, training their teams, and demanding that their technology partners bake compliance into the product rather than bolting it on after the fact.

And the technology companies that will endure aren't the ones making the boldest claims on their booth banners. They're the ones, like Anthropic, building deliberate friction into their products because they understand that in law, the cost of a mistake isn't a bad user experience—it's a \$10,000 sanction, a malpractice claim, or worse.

Agentic AI is coming for the legal profession. The only question is whether we'll be ready.

From Legalweek 2026, I can tell you: we're not. Not yet. But at least now we know what we're building toward.

Matthew A. Mishak, Esq. is the Founder & CEO of LegalTek.ai (d/b/a SilverTung), an AI-powered legal document automation platform built for Ohio domestic relations practitioners. He is also Managing Attorney of Mishak Law LLC in Amherst, Ohio, and Law Director for the Village of South Amherst. Follow his coverage of AI in the legal industry on [LinkedIn](#) and at [LegalTek.ai](#).

Legalweek New York 2026 runs March 9–12 at the North Javits Center. The "AI Agents 101: What They Are and Why They Matter" panel was part of the AI Agents and Automation educational track, supported by LexisNexis.