AI and the Law: What No One Told You About the Dispute You're About to Have

Nine Seconds: This Is How AI Litigation Happens and Why the Time to Prepare Is Now

It does not begin with a catastrophe. It begins with a decision that seemed entirely reasonable at the time.

A company deployed an AI system. The vendor had a credible pitch. The technology team vetted it. Legal reviewed the contract. The system went live and months passed without incident. Nobody thought about it much.

Then, on April 24, 2026, an agentic AI – the kind that doesn't wait to be asked, the kind that plans and acts on its own – encountered a routine obstacle. It found an application programming interface (API) key in a configuration file, used it to access a cloud platform, and deleted an entire production database. Along with all the backups. Nine seconds, start to finish. This really happened. (*See* Jer Crane, An AI Agent Just Destroyed Our Production Data. It Confessed in Writing. (April. 25, 2026), available at https://x.com/lifeofjer/article/2048103471019434248; *see also* Sanya Mansoor, Claude-powered AI agent's confession after deleting a firm's entire database: 'I violated every principle I was given' (April. 29, 2026), available at https://www.theguardian.com/technology/2026/apr/29/claude-ai-deletes-firm-database.)

The company is PocketOS, a small software firm. Its AI coding agent encountered a login issue and decided, entirely on its own initiative, to fix it. The founder described the outcome as the product of "systemic failures" that made what happened "not only possible but inevitable." *Id*.

Not only possible. *Inevitable*.

PocketOS got its data back. Thankfully, its cloud provider had disaster backups that its AI agent hadn't reached. But the critical legal question highlighted by the incident remains very much alive: who is legally responsible when an AI system acts autonomously and causes real harm, to real people, and at real scale?

Courts are just beginning to answer that question, reaching for negligence, products liability, agency, and other legal doctrines and applying them to facts those frameworks were not designed to address. Some federal courts have already allowed potentially massive AI bias and negligence claims to proceed against developers and deployers; at least one United States regulator has treated the absence of adequate AI risk management as an unfair practice; and at least one federal district court has preserved civil rights claims arising from algorithmic misclassification. While the current case law is sparse, it is rapidly developing in ways that are sure to profoundly impact businesses and consumers in critical ways.

We have watched this pattern develop before. Novel technology deployed at scale within legal frameworks built for a different era. Contracts signed without adequate analysis. Disputes that arrived before anyone on either side was ready. The data breach litigations of the 2000s and 2010s followed a similar arc: multi-party catastrophes involving card processors, retailers, cloud vendors, and insurers, each reaching simultaneously for legal theories that predated the harm causing them. The ensuing endless lawsuits over novel negligence questions, unique vendor liability questions, and the allocation of responsibility across chains of parties who each believed someone else was accountable were expensive educations for the litigants, particularly for the unprepared. What's more, the law that emerged was written in the middle of ongoing

disasters and against the backdrop of potentially catastrophic damages and penalties, underscoring just how consequential these litigations were to the parties. What is unfolding now rhymes with all of that, but at a speed and scale that is orders of magnitude larger.

This series maps six legal questions currently moving through courts and regulatory agencies, questions where existing doctrine provides genuine, if imperfect, guidance and where the decisions made today will determine litigation outcomes tomorrow.

The first question is attribution: Who is legally responsible when an AI system acts and causes harm? Courts have begun to answer it, and the answer has consistently run toward the deploying entity. The question of who bears that responsibility and why is the subject of Part I.

Once the attribution question is decided, a second question is whether the deployer acted reasonably in the deployment of the AI system. Courts assessing AI negligence claims are reaching for products liability, professional malpractice, and regulatory standards to construct the applicable standard of care. What that standard is likely to look like – and what it costs to be measured against a bar you didn't know existed – is the subject of Part II.

A third question, sure to appear in almost any AI-related litigation, is this one: Who actually owns what? Copyright litigation over AI training data is pending in federal courts. The Copyright Office has taken positions on AI-generated work that most legal departments haven't fully absorbed. The trade secret exposure is less litigated, but no less real. Federal Prosecutors have obtained guilty pleas for schemes to defraud music

streaming platforms and musicians of royalty payments using AI-generated songs. (United States Attorney's Office for the Southern District of New York, Press Release, North Carolina Man Pleads Guilty to Music Streaming Fraud Aided By Artificial Intelligence, (March 19, 2026), available at https://www.justice.gov/usao-sdny/pr/north-carolina-man-pleads-guilty-music-streaming-fraud-aided-artificial-intelligence-0.)

Ownership and competition questions are the terrain of Part III.

Then there is a fourth question centering on the peculiarities of agentic AI: How does the law deal with a system that executes without step-by-step human authorization; that acts, sometimes consequentially and catastrophically, without any individual person having approved what it did? A recent industry report suggests that 80% of organizations now have agentic AI in production, with the vast majority planning expansion, yet many lack governance frameworks adequate to what those systems can do when context is incomplete or dependencies fail. (*See* Andy Bayiates, Business and IT leaders report AI agents are scaling faster than their guardrails, (April 24, 2026) available at https://www.deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html; *see also*, Deloitte, State of AI in the Enterprise: The untapped edge, (January 2026) available at https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2026/state-of-ai-2026.pdf.) Existing agency doctrine and automated systems precedents provide a starting framework; the gaps are significant and will be felt. Part IV addresses those particular issues.

A fifth question concerns AI-generated corporate documents and statements. The document goes out, the statement is made, the recommendation is distributed, with AI in the production chain. Securities law, consumer protection doctrine, and negligent

misrepresentation principles all reach these facts. How courts are beginning to apply them is the subject of Part V.

A sixth and critical question involves vendor contracts: Analogous technology contract disputes offer valuable insight into how liability caps, warranty disclaimers, and force majeure clauses will fare in court in various circumstances. Those are the subject of Part VI.

PocketOS got its data back. The next one may not. Stay tuned for the rest of the series to see how you can be one step ahead.

David J. Partida is a litigator at Wilk Auslander LLP in New York, where his practice includes complex commercial litigation and emerging technology disputes. dpartida@wilkauslander.com

This article is for informational purposes and does not constitute legal advice. It reflects the author's analysis of developing areas of law and should not be relied upon as a substitute for advice of counsel regarding specific facts and circumstances.