ComputerForensic Services

NOVEMBER 2026

PERSPECTIVE · TECHNOLOGY & THE LAW

Autonomous AI and Digital Evidence

When Nobody Pressed "Send"

By Mark Lanterman



As autonomous AI systems begin taking actions on behalf of users, traditional questions of attribution, authenticity, preservation, and digital evidence become significantly more complex.

Visualize this situation. A law practice decides to adopt an artificial intelligence assistant. Lawyers use the technology to write emails, set appointments, organize files, and talk to clients. One morning, privileged client material turns up in opposing counsel's inbox.

The firm's managing partner wastes no time asking a familiar question:

"Who sent that email?"

It's not an easy question to answer.

No lawyer wrote the message. No paralegal hit "send." In fact, no person pressed a single key. The email was composed and sent by an autonomous AI agent acting upon instructions it had been given earlier.

Say hello to the future of digital evidence.

Most conversations involving generative AI have centered on deepfakes, fabricated content, and AI driven misinformation. Although problematic, these issues generally have one thing in common. A human being instructed the technology to create something.

Not so with the next generation of AI.

Instead of passively generating information, autonomous AI agents will be instructed to take action. These tools can set appointments, retrieve records, prepare documents, talk to customers, initiate transactions, and complete a wide variety of tasks without human intervention. Imagine technology companies' vision of the future coming true. Lines of code operating 24 hours a day as digital workers performing tasks on behalf of companies.

Attorneys will have a host of new questions.

Attributing digital evidence to a specific person has been a major focus of computer forensic investigations for years. Did someone open an email? Did someone access a database? Did someone approve a transaction? Forensic examiners have hunted down the human perpetrators by reviewing metadata, system logs, cloud storage records, and artifacts left on electronic devices. All that happens after a person takes action.

> *"No person pressed a single key."*

What happens when AI software takes the action?

Consider the following.



An attorney instructs an AI software agent to "file the quarterly report." The agent accesses the necessary records, prepares the report, uploads the file to the relevant platform, and sends the report to various stakeholders. If something goes wrong with the report, who did it? The lawyer who instructed the agent? The developer who wrote the software? The company? The AI system?

Accountability is only part of the problem.

Questions of authenticity will also arise. For decades, courts have asked whether a document is what its proponent says it is. To answer that question, lawyers have traditionally identified the person who created or sent the evidence. When AI agents act on their own, will lawyers have to identify which software program took an action? And even after you find a responsible party, intent could be the most difficult thing to prove, because a rogue agent can behave in unintended or unforeseen ways.



And what about discovery?

Expect litigants to request production of AI prompts, agent instructions, automation histories, system logs, and AI model outputs. Will courts require companies to preserve records that explain how an AI software agent decided to take action? What about when those records are automatically overwritten, deleted, or not created in the first place?



Autonomous AI will also raise novel spoliation concerns.

Everyone understands that a party must preserve evidence when litigation is pending or reasonably foreseeable. But AI systems can create massive amounts of information. Not only will it be difficult to determine what needs to be preserved, but it will be equally challenging to determine when the duty to preserve attaches.

⚠ REAL-WORLD WARNING SIGNS — INCLUDING IN TENNESSEE

Sounds like a future case? It is already here. In December 2025 and again in March 2026, the Tennessee Board of Professional Responsibility publicly censured lawyers who filed briefs containing fictitious citations generated by AI tools and never verified. In the second matter, the lawyer believed a paralegal had checked the citations but never confirmed it—and was censured under RPC 5.3 for failing to supervise. That is the "who did what?" problem in miniature, and it arose with a human still in the loop. As AI systems move from drafting to acting—talking to clients, completing transactions, filing documents with little human intervention—the toughest question will not be whether an AI system did something, but whether anyone can show why it did it, and who was supposed to be watching.

Things have not gone uncontested in the past.

Email wasn't discovered until someone figured out how to review email. Mobile devices weren't subject to search warrants until attorneys identified what was actually on people's phones. Cloud computing caused a preservation crisis—until courts developed workable solutions. Social media sparked questions about authenticity and credibility—until lawyers learned how to investigate social media evidence.

It's only a matter of time before autonomous AI follows the same path.

The better news is that the law probably won't change.

Courts will likely continue to hold parties responsible for the authenticity, reliability, and credibility of the information they introduce. What will change are the technologies that produce evidence. Software doesn't testify in court, but that doesn't mean lawyers should ignore the records related to how AI software is instructed to act.

Attorneys should start preparing today.

The professional rules already supply the framework. Tennessee RPC 1.1 (and its comment on keeping abreast of relevant technology), RPC 1.6 (confidentiality), and RPC 5.3 (responsibility for nonlawyer assistance) apply whether the "assistant" is a paralegal or a software agent—a point ABA Formal Opinion 512 (2024) makes explicit for generative AI. Firms and businesses that deploy autonomous AI should adopt policies on acceptable use, audit trails, record retention, and human supervision. Did someone instruct the AI? Did someone review the AI's actions? The answers to those questions (and the records that support them) may one day be critical to litigation—or to a disciplinary inquiry.

SOURCES

Bd. of Prof'l Resp. of the Sup. Ct. of Tenn., Public Censure of Matthew B. Reeves, BPR No. 035871 (Dec. 15, 2025); Public Censure of Matthew W. Willis, BPR No. 024803 (Mar. 30, 2026), both available at tbpr.org.

Tenn. Sup. Ct. R. 8, RPC 1.1 & cmt. [8]; RPC 1.6; RPC 5.3. · ABA Standing Comm. on Ethics & Prof'l Resp., Formal Op. 512, *Generative Artificial Intelligence Tools* (July 29, 2024).

For decades, digital forensic examiners have asked:

WHO DID WHAT?

Soon they may need to ask:

DID ANYONE DO ANYTHING AT ALL?

That question may define the future of digital evidence.



ABOUT THE AUTHOR

Mark Lanterman

CTO, COMPUTER FORENSIC SERVICES

Nationally recognized for his work in digital forensics, cybersecurity, and electronic evidence, Mark has more than 30 years of experience assisting attorneys, corporations, and government agencies with complex technology investigations. A former member of the U.S. Secret Service Electronic Crimes Task Force, he has testified in more than 2,000 cases and has trained judges, attorneys, and investigators throughout the country.

ComputerForensic Services

Computer Forensic Services
800 Hennepin Avenue, Fifth Floor
Minneapolis, MN 55403
952-924-9920 · www.compforensics.com