

New York Skipped the AI Disclosure Fight. It Went Straight to Human Accountability.

**Part 161 is not a permission slip for lawyers to use AI quietly.
It is a checkpoint rule in legal form.**

**Basil C. Puglisi, MPA
A Human-AI Collaboration**

New York did not make AI disclosure the center of legal accountability. Instead, they made it about the human checkpoint, demanding a signature and professional accountability.

When Part 161 took effect on June 1, 2026, it did not require attorneys or parties to file a routine disclosure form every time artificial intelligence helped prepare a court paper. That choice matters. The rule permits AI-assisted drafting, research, editing, summarizing, and organization, but it does not move responsibility from the lawyer to the tool. It keeps accountability attached to the person who submits the work.

That is the real governance move.

Part 161 does not treat AI use as the legal event. The filing is the legal event, and the signature is the certification point. If an AI-assisted submission contains fabricated cases, fictitious statutes, false factual statements, or frivolous legal arguments, the absence of a disclosure requirement does not protect the signer. The lawyer or party still owns the paper, still owns the verification duty, and still answers for the result.

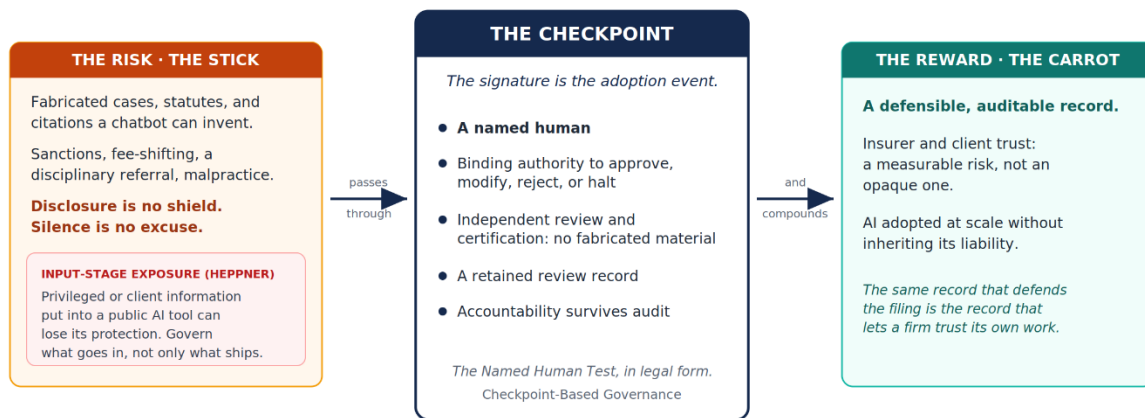
This is why Part 161 should not be read as permission for casual AI use. It is better understood as a shift from disclosure-based governance to checkpoint-based accountability. The court does not need every filer to announce the software. It needs every filer to verify the submission before placing it into the record.

That approach did not appear in isolation. In October 2025, the New York State Unified Court System adopted an interim AI policy for judges and court staff that allowed limited AI use while preserving human judgment, review, confidentiality, and accountability. Part 161 carries that same logic into attorney and party filings. AI may assist the work, but it does not absorb the responsibility.

The named human still does.

Part 161: The Risk, the Checkpoint, and the Reward

New York moved the governance event from disclosing the tool to owning the output.



The court does not ask which tool was used. It asks who signed.

Basil C. Puglisi, MPA · A Human-AI Collaboration · #AIassisted using the HAIA Ecosystem

What Part 161 actually says

Part 161 was added by Administrative Order AO/75/26, signed by Chief Administrative Judge Joseph A. Zayas on March 25, 2026, and took effect on June 1, 2026. It applies to all courts of the Unified Court System in both civil and criminal cases. It defines AI broadly, as a machine-based system that makes predictions, recommendations, or decisions for human-defined objectives using model inference, and it defines a paper as a brief, memorandum, affidavit, affirmation, pleading, or other document prepared for submission to a court. It excludes materials offered as evidence, which raise separate questions the rule leaves alone.

Attorneys and parties should not be prohibited from using AI tools to prepare papers, as long as that use complies with the duties already attached to court submissions. Because those duties apply whether or not AI was used, the rule concludes that attorneys should not be required, simply upon filing, to disclose that AI helped prepare the paper. New York does not treat AI use as the event that matters most. It treats submission to the court as the event that matters most.

What the model rule in Appendix A adds

If the body of Part 161 sets the policy, the teeth sit in Appendix A, a model rule that individual courts and judges are encouraged, but not required, to adopt. Where a court adopts it, the model rule expects every attorney or party who uses an AI tool to understand that tool's capabilities and limitations. It warns, in the rule's own language, that AI tools can generate fabricated information and fictitious citations, the failures

commonly called hallucinations. It then ties that warning to existing authority. By signing and submitting a paper, an attorney or party certifies, after reasonable inquiry, that it contains no false material factual statement and no frivolous legal argument, a certification rooted in 22 NYCRR 130-1.1 and 130-1.1a, and the attorney remains bound by Rule 3.3 of the New York Rules of Professional Conduct.

Where the model rule applies, the signer must carefully review the paper and independently ensure that it contains no fabricated or fictitious cases, statutes, or other material, and by signing certifies that this review was done. If a court finds that the requirement was not met, it may impose sanctions or other remedial action. The certification duty is not new; it predates AI, and the model rule simply points that duty at a new source of error. Adoption is also left to individual courts, so a New York practitioner has to know which courts have adopted the model rule and which judges have added their own AI requirements, since several already have.

No disclosure required, total accountability retained

The absence of a disclosure rule is a trap rather than a relief. New York permits AI use without routine disclosure, but that does not make the use safer. It makes the review more important. If a tool is used quietly and the filing is wrong publicly, the signature is still the checkpoint. The court does not need to know which tool produced the error in order to ask who filed it. The attorney did.

Saying AI helped does not cure a fabricated citation, and not saying AI helped does not excuse one. The duty attaches to the filing, and the consequences attach to the signer. A fabricated case in a brief can draw sanctions, fee-shifting, a disciplinary referral, a stricken filing, and malpractice exposure that follows the lawyer long after the case ends. None of that turns on whether AI was disclosed. All of it turns on whether the signer ran a real review before the paper became official.

In *Mata v. Avianca*, the Southern District of New York sanctioned lawyers who filed a brief containing fabricated cases generated by a chatbot, cases they had not verified. The deeper failure was not that the tool produced fiction. It was that the fiction moved from machine output to court filing without adequate human verification, and the human checkpoint gave it institutional force. Part 161 responds to that same failure pattern. It does not try to ban the tool or shame the use, and it does not make disclosure the remedy. It returns the court's attention to the professional act that has always mattered, the signed filing. *Mata* arose in federal court, and Part 161 governs New York's state courts, so the two are not the same forum. The duty of candor and the verification expectation, though, run across both forums, which is why the state rule reads as a response to a federal lesson.

New York lawyers carry a second AI exposure that Part 161 does not address, and honesty about the landscape requires naming it. The risk is not only at the filing checkpoint. It can begin the moment information goes into the tool. In *United States v. Heppner*, the Southern District of New York held that a defendant's self-initiated documents, created by querying a publicly available AI platform and later shared with counsel, were protected by neither attorney-client privilege nor the work product doctrine, because the platform was not a lawyer, its terms allowed the provider to use and share inputs, and the materials were not prepared at counsel's direction. *Heppner* is a confidentiality ruling, not a fabrication ruling, so it sits at a different stage than Part 161. The lesson points the same direction one step earlier. Putting privileged or client information into a public, consumer-grade AI tool can jeopardize the protection a firm assumes it has, which means a real internal checkpoint has to govern what goes into the tool as carefully as what comes out of it. The holding is fact-bound, and the court suggested that a tool used at counsel's direction might be treated differently, so a secured, enterprise tool used under a lawyer's direction is not the same case as a defendant's solo use of a public chatbot.

The checkpoint that protects you is the one that scales you

Whatever else governance buys, it does not yet command a clear price in the market. The recognized currency today is the defensible record itself, not a premium credit or a client's stated preference. Any claim that governance already pays in dollars would get ahead of the evidence, so what follows is an argument about where the incentive points, not a report on a market that has already arrived.

The case for a real internal checkpoint usually stops at the fear of sanctions, and that is incomplete because the same work produces an upside. The first reward is a record. The citation checks against authoritative databases, the comparison of quotations to the original authority, the named attorney's sign-off, and the retained evidence that the review happened are not bureaucracy. They are the materials that answer a sanctions motion, give a malpractice carrier something measurable to weigh, and reassure sophisticated clients who increasingly ask not whether a firm uses AI, but how its AI-assisted work is verified before it becomes legal action. The firm that builds that record early is not merely performing diligence. It is assembling proof a careful field is starting to expect, while the proof is still cheap to assemble.

Underwriting logic points the same way, though no carrier has yet promised a price for it. Insurers assess what they can measure: recognized controls, attestations, and loss history. A firm that can show a documented AI verification process presents a measurable risk; a firm that cannot presents an opaque one, and an opaque risk is the kind that tends to draw more questions rather than fewer. That is the shape of the incentive, not a premium outcome on any current rate sheet. It is the point the author

has argued elsewhere: insurers cannot price what governance cannot prove, and the signed, reviewed, documented filing is the kind of proof a carrier would one day need. The same logic reaches clients and courts, who extend trust more readily to counsel who can show their work than to counsel who ask to be taken on faith.

One could say the signature has always carried this duty, so Part 161 changes nothing and the reward is just doing the job. That is half right. The duty is old, but the AI era multiplies both the volume and the plausibility of fabricated material, and a confident, well-formatted false citation is now cheap to produce at scale. The firm that builds its verification record now assembles it before it is tested in a sanctions hearing, not after, which is the only time the record is affordable to build. A second objection runs the other way: permitting AI without disclosure invites quiet misuse. Disclosure never cured a bad filing, and its absence does not excuse one. The accountability rule reaches the bad filing regardless of whether the tool was ever named.

This is why the defensive case and the growth case converge on the same work. The discipline that produces the record, deciding who is accountable, checking the sources, and logging where a human exercised judgment, is also the discipline that lets a firm trust its own AI-assisted output enough to use it broadly. Governance is the competitive asset. The tool is available to every firm; the discipline to govern it is not. The firm that governs the work well is not choosing between defending itself and growing. It is doing both with the same work.

The signature is both shield and engine

The named human at the filing checkpoint, holding binding authority and leaving a retained record, is at once what answers a sanctions motion and what lets a firm rely on its own AI-assisted work, which is why the two rewards arrive together. The same act does both jobs. That is not a coincidence of the rule. It is the structure of governance itself, expressed in the one place every filing already passes through.

The author has called this Checkpoint-Based Governance, and Part 161 fits that model at the filing checkpoint. Real governance does not begin when a person is somewhere near the system. It begins when a named human holds binding authority at a defined checkpoint and can approve, modify, or reject the output before it becomes consequential. The Named Human Test asks a single question of any AI-assisted process: is there a named human who holds binding authority to approve, modify, or reject the output, and whose accountability survives audit? If the answer is no, the process may be careful, automated, and well documented, but it is not governed. If the answer is yes, governance can begin, because authority and consequence attach to a person. A lawyer's signature is one of the oldest answers to that test in professional life. It is not decorative, and it is not a workflow marker. It is a formal act of adoption, and Part 161 modernizes it for the AI era.

This is also where New York's choice connects to the wider standard of care without being swallowed by it. The accountability-first posture of Part 161 runs in the same direction as the NIST AI Risk Management Framework and ISO/IEC 42001, the governance baselines that, as a standard of care forms, a careful field should be expected to treat as the benchmark for reasonable AI conduct, a development the author has examined elsewhere. It runs on a different logic than the European Union's approach, where the General Data Protection Regulation gives individuals rights against decisions made solely by automated processing that carry significant effects, and the AI Act's Article 50 requires certain AI-generated content to be disclosed and labeled, both leaning on rights and disclosure rather than a signer's certification. New York did not adopt either model. It reached for the instrument already in every lawyer's hand, the signature, and made that the point where AI-assisted work becomes a human's responsibility, which is the difference between a human in the loop and a checkpoint. A loop is a diagram. A checkpoint is a decision.

None of this is automatic

A rule that locates governance in the signature only works if the signature means something. A named attorney who lacks the time, the competence, the access to source materials, or the authority to change or stop a filing is a name, not a checkpoint, and a signature given without independent review becomes evidence against the process rather than protection for it. Human in the loop is not enough because a person can be present and still hold no authority and bear no real accountability. The model rule rewards the function the signature performs, not the formality of signing it.

The record can also cut the wrong way. A verification process performed badly, or a review log built carelessly, can be discovered later and used to show that a firm claimed a standard it did not keep. A standard adopted is not a standard honored. The protection Part 161 offers is real, but it is conditional on the review behind the signature being real, and that condition is the whole of it.

Built to outlast the tools

The reason to settle this now is that AI is ceasing to be a separate tool a lawyer chooses to open and becoming part of the software a law office already runs. AI features are being built into word processors, legal research platforms, transcript tools, discovery systems, document management, and e-filing software, and in that environment a disclosure-first rule becomes both overbroad and underinclusive at once. It can force a lawyer to disclose trivial assistance while missing AI functions buried inside ordinary software the lawyer never thinks of as AI. An accountability-first rule survives that shift because it does not depend on the lawyer identifying every embedded feature in the stack. It depends on whether the lawyer can certify the paper.

That also makes Part 161 more durable as legal work becomes agentic, and tools move from drafting help to research sequencing, document comparison, transcript summarization, and multi-step litigation support. The disclosure question gets less useful as the tools multiply. The better question holds steady: did a competent lawyer review and certify the filing before it entered the court record? New York has not said, "Use AI and tell no one." It has said something sharper. Use AI if you choose, but when you sign, the output is yours.

Sources and currency

New York State Unified Court System. (2026). *Administrative Order of the Chief Administrative Judge AO/75/26* (Mar. 25, 2026), adding a new Part 161 (22 NYCRR 161.1 to 161.4 and Appendix A) to the Rules of the Chief Administrator of the Courts, effective June 1, 2026. <https://www.nycourts.gov/rules/part-161-use-artificial-intelligence-technology>

New York Codes, Rules and Regulations. *22 NYCRR 130-1.1 and 130-1.1a* (frivolous conduct and sanctions; signing of papers and certification after reasonable inquiry).

New York Rules of Professional Conduct, Rule 3.3 (22 NYCRR Part 1200) (candor toward the tribunal).

Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023).

United States v. Heppner, No. 25-cr-00503-JSR, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026).

New York State Bar Association. (2026). *Effective June 1, 2026, the New York State Unified Court System has adopted a new rule regarding the use of artificial intelligence.* <https://nysba.org/effective-june-1-2026-the-new-york-state-unified-court-system-has-adopted-a-new-rule-regarding-the-use-of-artificial-intelligence/>

New York State Unified Court System, Office of Court Administration. (2025, October 10). *Interim Policy on the Use of Artificial Intelligence* (press release PR25_23). https://www.nycourts.gov/LegacyPDFS/press/pdfs/PR25_23.pdf

Legal AI Governance Tracker. (2026). *New York* (individual-judge AI rules and Part 161 adoption status). <https://legalaigovernance.com/tracker/states/new-york/>

National Institute of Standards and Technology. (2023). *AI Risk Management Framework (AI RMF 1.0, NIST AI 100-1)*. <https://www.nist.gov/itl/ai-risk-management-framework>

International Organization for Standardization. (2023). *ISO/IEC 42001:2023, Artificial intelligence management system*.

European Union. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation)*, Article 22 (automated individual decision-making).

European Union. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*, Article 50 (transparency obligations for certain AI systems). <https://artificialintelligenceact.eu/article/50/>

Puglisi, B. C. (2026). *The Standard of Care: How NIST and ISO Are Turning Voluntary AI Governance Into a Liability Defense*. <https://basilpuglisi.com/standard-of-care-ai-governance/>

Puglisi, B. C. *Checkpoint-Based Governance*. Author-proposed operating model. <https://basilpuglisi.com/checkpoint-based-governance/>

Puglisi, B. C. (2026). *The AI Risk Economy*. <https://basilpuglisi.com/ai-risk-economy-insurance-governance/>

Disclaimer. The author is not a lawyer, and nothing in this article is legal advice. Part 161, the cases discussed here, and the duties they describe should be confirmed against the current rule text and authority, and any decision about AI use in legal practice should be made with qualified counsel. The author is an independent practitioner and author who may profit in other ways from research and content like this.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance

[Spotify](#); [Apple Podcasts](#); [Amazon Music](#); [YouTube Playlist](#);

#AIassisted using the HAIA Ecosystem