

# Guardrails Before Greenlights: How Gen AI Will Actually Shape E-discovery in 2026



JANUARY 6, 2026

*This article was originally published in [LegalTech News](#). Any opinions in this article are not those of Winston & Strawn or its clients. The opinions in this article are the authors' opinions only.*

Generative AI will be everywhere in e-discovery by 2026, but not necessarily where some evangelists predict. Call us contrarian if you must, but we do not believe the headline-ready notion that large language models will substantially supplant traditional responsiveness review for production will materialize this year.

Instead, 2026 will see generative AI become integral across a host of document-related litigation workflows, including issues analysis, privilege screening, quality control, sensitivity detection, production analysis, chronology building, fact development and deposition prep, while human judgment remains central to responsiveness determinations for outgoing productions. At the same time, we expect courts finally will begin to articulate the contours of defensible AI use in production workflows, offering early guardrails without endorsing wholesale automation of responsiveness calls.

## THE LEGAL POSTURE: NO “DA SILVA MOORE” MOMENT—YET

The two predictions are interrelated. Simply put, many sophisticated clients in high-stakes litigation are—and will remain—wary about using generative AI as a primary decision-maker in a responsiveness review workflow until courts have had their say. As of the close of 2025, the judiciary has not delivered the kind of watershed opinion for generative AI that Judge Andrew Peck’s decisions in *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012) (“computer-assisted review is an available tool and should be seriously considered for use in large-data-volume cases where it may save the producing party (or both parties) significant amounts of legal fees in document review”), and *Rio Tinto PLC v. Vale S.A.*, 306 F.R.D. 125 (S.D.N.Y. 2015), provided for technology-assisted review over a decade ago.

There currently is no authority declaring the use of generative AI for responsiveness determinations to be reasonable. The few opinions that have touched the edges of the issue—addressing counsel’s (and experts) duty to verify AI outputs, the adequacy of “AI agent” search approaches or cautionary tales about hallucinated citations—are instructive but not e-discovery-specific and, critically, not broadly transferrable to corporate defendants’ production workflows.

This unsettled landscape matters. Without clear standards for defensibility, most corporate litigants and their outside counsel will remain cautious about deploying generative AI on the front lines of outgoing productions. Courts have yet to weigh in on core questions, including whether generative AI may be relied upon to make final responsiveness determinations, how error rates should be measured and reported, what validation protocols are sufficient, and where and how human oversight must be documented. Until authoritative guidance and clear standards emerge, the risk-reward calculus will favor caution in 2026.

### **WHAT WILL CHANGE IN 2026: THE FIRST FRAMEWORKS EMERGE**

Even if 2026 will not normalize AI-reliance on responsiveness review, it is likely to mark the beginning of a jurisprudence foundation. Expect a first wave of opinions addressing defensibility for e-discovery generative-AI review—for example, in areas such as validation and sampling protocols; transparency expectations for meet-and-confers; disclosure obligations when Gen AI materially drives production or culling decisions; disclosure and discoverability of prompts; and parameters for auditing model outputs. Courts are likely to treat generative AI by analogy to prior TAR jurisprudence, emphasizing reasonableness, proportionality, cooperation and iterative validation while acknowledging risks unique to generative systems, including hallucination, model drift and explainability gaps.

Those touchstone opinions will not immediately greenlight AI as a replacement for human review for outgoing production decisions, but they will begin to define the practices that may eventually support that step. Think of 2026 as the year the scaffolding goes up: not yet a finished structure, but enough to begin building repeatable, defensible workflows. However, since we don't yet have our Da Silva Moore for generative AI, there is a distinct risk profile and a corresponding need for heightened, intentional focus on documented defensibility controls if litigants contemplate leveraging generative AI for outgoing production decisions.

Because this remains largely uncharted territory, defensibility will hinge on disciplined implementation. Litigation teams should calibrate generative AI to tasks that tolerate probabilistic outputs, build human-in-the-loop checkpoints, and document their process and results. They should also approach AI use strategically and constructively during meet-and-confers, where the facts and circumstances warrant such discussions and transparency.

### **THE PRODUCTION LINE WILL WAIT, BUT NOT FOREVER**

The practical upshot for 2026 is straightforward. Do not expect a broad, industry-wide shift to generative AI as the primary engine for responsiveness and privilege determinations in document productions. That step will likely require clearer standards, shared benchmarks, and a handful of influential opinions blessing specific methods. We are not there yet, but we are knocking on the door.

What you should expect is nearly universal adoption of generative AI for efficiency-driven tasks that accelerate understanding without assuming the burden of certifying outputs to courts and parties. Teams that embrace these tools thoughtfully—focusing on evidence storytelling, sharpening and shortening preparation, and extracting insights from adversary productions—will gain meaningful strategic advantages while staying on solid defensibility footing.

### **WHERE GEN AI WILL SHINE IN 2026**

In the meantime, generative AI will transform the handling of litigation data in ways that are powerful, measurable and arguably lower-risk. Its strengths are especially well-suited to high-volume, interpretive, and narrative tasks that stop short of making final production calls. By the end of 2026, expect routine use across workflows such as document summarization at scale, chronology building and timeline creation that integrate custodial collections, deposition preparation with witness-specific briefing books surfacing key documents, themes, and likely lines of questioning while rapidly comparing inconsistent statements across datasets; analysis of incoming and third-party productions (where the defensive posture reduces risk and the goal is understanding and locating key documents rather than certifying the reasonableness of the process); and issues-spotting and theme alignment that help trial teams connect documents to legal theories, jury narratives, and expert frameworks without supplanting human conclusions.

These uses offer material efficiency gains while avoiding the heightened defensibility burden that accompanies outgoing productions. Properly deployed, they reduce cycle time and cost, improve consistency, and surface insights

earlier—advantages that compound across matters and portfolios. And for e-discovery and litigation teams that are properly trained and equipped, they are available and accessible right now.

## BOTTOM LINE

In 2026, generative AI will not replace human judgment for responsiveness reviews in outgoing productions; it will augment it everywhere else. The year will be a turning point, not because courts will declare generative AI per se reasonable for responsiveness, but because they will begin to articulate the rules of the road.

Until a true “Da Silva Moore moment” arrives for generative AI in e-discovery, the winning strategy is prudent adoption: Use it broadly for insight and speed, back it with rigorous validation and human oversight and keep the high-stakes calls in workflows you can defend. Those who strike that balance will move faster, see clearer and avoid avoidable risk—precisely the edge that matters in modern litigation.

5 Min Read

---

## Related Capabilities

[eDiscovery & Information Governance](#)

[Artificial Intelligence \(AI\)](#)

## Related Professionals

---



[Bobby Malhotra](#)



[Jason Moore](#)