

Court Evaluates Generative AI Review Under Established Discovery Principles

By Daniel J. Toal, Geoffrey R. Chepiga and Ross M. Gotler


August 3, 2026

In *Schulte v. LinkedIn Corp.*, 2026 WL 1905851 (N.D. Cal. July 1, 2026), Magistrate Judge Laurel Beeler relied on two settled principles of discovery law to resolve a dispute over a defendant's use of a generative artificial intelligence (AI) tool to decide which documents were responsive.

First, a responding party is generally best positioned to determine how to search and review its own data, and that a court will not second-guess those choices absent a concrete showing that the process has failed. Second, a party seeking "discovery on discovery" must identify a specific deficiency in the production, not merely speculate that responsive material has been missed.

But what makes the decision most notable is what it does not do. The court declined to treat the defendant's use of that AI tool as anything out of the ordinary. It did not declare that it was approving the use of AI, and it did not subject the tool to additional scrutiny. Rather, the court analyzed the tool as it would any review method and applied well-established standards of relevance and proportionality.

That restraint is itself telling. In 2012, when many lawyers doubted that courts would accept technology-assisted review (TAR), then-Magistrate Judge Andrew J. Peck wrote the first judicial



Credit: Yingyapumi/Adobe Stock

opinion to approve it, reassuring the bar that counsel "no longer have to worry about being the 'first' or 'guinea pig' for judicial acceptance of computer-assisted review." *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182, 193 (S.D.N.Y. 2012).

The *Schulte* court needed no such pronouncement, applying the rules that have long governed keyword searching and TAR to generative AI review, and requiring no special approval or heightened scrutiny.

Generative AI Document Review

The plaintiffs in *Schulte* are a putative class of LinkedIn Premium subscribers who allege that LinkedIn monopolized a market in violation of Section 2 of the Sherman Act. The central dispute before the court concerned defendant LinkedIn's search-and-review methodology.

As discovery opened, LinkedIn gave plaintiffs twenty-five search strings that it intended to run against its custodial data. LinkedIn also disclosed that after running those searches it would use Relativity aiR, a generative AI review platform, to review the resulting hits for responsiveness.

When plaintiffs asked for more information, LinkedIn “disclosed that ‘(1) there was no seed or training set used; (2) Relativity aiR is being used to make final responsiveness calls; and (3) quality control review is being conducted by human review of samples taken from each responsiveness type.’”

Plaintiffs objected, taking the position that LinkedIn’s disclosures were “inadequate” and its review process was “inappropriate.” After the parties were unable to resolve the dispute, plaintiffs moved the court (i) to bar LinkedIn from applying search strings to pre-cull the documents before using Relativity aiR, (ii) to require LinkedIn to run aiR across every custodial file the parties had agreed on, and, separately, (iii) to compel LinkedIn to disclose additional metrics detailing its use of aiR.

Search-Term Pre-Culling

Arguing against the use of search strings prior to generative AI review, plaintiffs contended that such culling was unnecessary and that it “artificially reduce[s]” the population of documents the tool would see, screening out responsive material before the tool reviewed it.

LinkedIn responded that its culling was proper, that it had met and exceeded its disclosure obligations under the parties’ Interim ESI Order, and that applying the tool to all custodial files without any culling would be disproportionate and impose an undue burden on it.

The court declined to bar pre-review culling. It observed that plaintiffs had never contended the twenty-five search strings were too narrow, nor voiced any objection to them when the defendant first produced the list. See *id.* Settled

authority agreed, as other courts “have held that using search terms to pre-cull documents before providing them to technology-review platforms satisfies the reasonableness and proportionality standards of Rules 26(b) and 34(b)(2).”

The court explained that if plaintiffs believed the strings were too narrow and risked leaving responsive documents unproduced, the proper course was to press for broader or revised search strings rather than to bar culling altogether.

The court also found that barring pre-review culling entirely “would impose a substantial burden on LinkedIn.” Without pre-review culling, LinkedIn would have to run multiple terabytes of material through the platform for all nineteen custodians (the files of just two custodians already came to roughly 800 gigabytes), at “significant costs related to processing, hosting, and human review.” The court directed the parties to meet and confer about the search strings within twenty-one days, with leave to return if they could not agree.

Relativity aiR Metrics

Plaintiffs additionally moved to compel the defendant to produce a set of performance metrics for its Relativity aiR review, “including elusion estimates, the document error rate, and the number of human reviewers being used to validate Relativity aiR’s predictions.” This request fared no better. The court treated the request as a form of “discovery on discovery,” an attempt to probe into the defendant’s discovery process.

Quoting a recent decision from the same district, it explained that discovery on discovery “is disfavored, as such discovery is typically not relevant to the merits of a claim or defense, and is rarely proportional to the needs of a case,” and is available only on a showing of a “specific deficiency” rather than “mere speculation.”

The court denied the request on two grounds. First, LinkedIn had already satisfied the Interim ESI Order, which only required a producing party

to disclose an intent to use technology-assisted review to screen out non-responsive documents. LinkedIn had done exactly that by identifying Relativity aiR, “a form of Technology Assisted Review,” and provided more detail when the plaintiffs asked.

Second, plaintiffs offered nothing more than the size of the target population, some 204,444 documents, as their claimed deficiency. But absent some concrete indication of a specific deficiency, “the number of documents in the target population alone does not warrant the requested discovery on discovery.”

Takeaways

Schulte returns repeatedly to a single theme: a generative AI review platform is a discovery tool, and a party that seeks to constrain an opponent’s process or audit its results must make the same concrete, particularized showing that has long been required, whatever the review technology. Several practical lessons follow.

First, a party that adopts generative AI for document review should, as with any review method, build a reasonable process and be ready to stand behind it. What made LinkedIn’s use of Relativity aiR defensible was not the tool itself but the process built around it, including sampling and validation.

The court treated generative AI review as a form of technology-assisted review and measured it against the familiar metrics of reasonableness and proportionality, not a heightened standard.

Second, a court’s deference to a responding party’s methodology is at its strongest when that methodology is facially sound. Unless constrained by restrictive ESI protocol language, a party is generally free to choose reasonable methods for meeting its discovery obligations,

including a decision to cull a set of documents with search terms before review.

When that party can stand behind the specific choices it has made, it is well positioned to defeat challenges to its process without opening its internal methodology to inspection merely because an opposing party takes issue with the approach.

Finally, requests for metrics or for an audit of an opponent’s review will generally be viewed as disfavored discovery-on-discovery. A party seeking such information typically must first identify a specific, concrete deficiency in the production, supported by evidence that responsive material is actually missing. As in *Schulte*, neither the size of a review population nor general unease about a process will justify such demands.

Generative AI has moved quickly into the mainstream of discovery, and parties are increasingly relying on it in the document review process. *Schulte* reflects that arrival.

The court saw no need to authorize the technology, to subject it to special scrutiny, or to require the producing party to justify its use. Just as with keyword searches or more established forms of technology-assisted review, the question is not whether a generative AI review tool is allowed, but whether the process around it is sound.

For federal practitioners who may have hesitated to evaluate or adopt generative AI review tools for fear of judicial second-guessing, *Schulte* offers reassurance: a well-designed and validated generative AI review process stands on the same settled footing as the review methods that came before it.

Daniel J. Toal and **Geoffrey R. Chepiga** are litigation partners and **Ross M. Gotler** is deputy chair and counsel, e-discovery at Paul, Weiss, Rifkind, Wharton & Garrison. **Lidia M. Kekis**, e-discovery attorney, assisted in the preparation of this article.