

Discovery Ruling Shows How Extant Rules Mesh With AI Tools

By **Christopher Davies, Marissa Wenzel and Zviad Guruli** (July 28, 2026)

As parties continue to employ generative artificial intelligence in discovery, a recent decision from U.S. Magistrate Judge Laurel Beeler in the U.S. District Court for the Northern District of California provides welcome guidance on the legal framework courts will apply to disputes about the use of generative AI tools in document review.

In *Schulte v. LinkedIn Corp.*, LinkedIn disclosed that it planned to use Relativity aiR, a generative AI-powered document review product, to assist in reviewing custodial documents for responsiveness. The plaintiffs raised several challenges to LinkedIn's plan, but on June 30, the court rejected all of them.

The underlying litigation is a putative class action bringing monopolization claims under Section 2 of the Sherman Act. Alleging two categories of anticompetitive conduct, the plaintiffs assert that LinkedIn overcharged its premium subscribers.

The alleged misconduct consisted of: (1) offering LinkedIn's private user data to potential rivals on the condition that the rivals not compete with LinkedIn; and (2) integrating user data with the Azure cloud product belonging to LinkedIn's parent company Microsoft, thereby driving up prices for computing resources.

Under the parties' protocol for electronically stored information, a producing party was required to disclose whether it intended to use technology-assisted review, or TAR. Consistent with that requirement, LinkedIn informed the plaintiffs that it intended to use Relativity aiR to filter out nonresponsive documents. Additionally, LinkedIn shared that it planned to use 25 search strings to "pre-cull" nonresponsive documents out of the set that would be fed into Relativity aiR.

In response to the plaintiffs' request for more information, LinkedIn also disclosed that it intended to deploy Relativity aiR to make final responsiveness calls with no seed or training set, and confirmed that its quality control would include human sampling of each responsiveness category.

Generative AI review is treated as TAR.

The most consequential aspect of the decision is that the court treated generative AI review as familiar rather than as a novelty. Instead of creating a separate legal framework for generative AI technology, the court analyzed generative AI review under principles developed in connection with traditional TAR.

By grounding its analysis in established TAR frameworks, the court signaled that existing proportionality, reasonableness and transparency standards can accommodate generative AI workflows without requiring a wholesale rethinking of discovery doctrine.



Christopher Davies



Marissa Wenzel



Zviad Guruli

This framing reflects the right analytical approach. Generative AI technologies should generally be treated as the "technology" in TAR.

As illustrated in the second edition of the Sedona Conference TAR Case Law Primer, litigants have leveraged artificial intelligence tools for well over a decade under the same TAR umbrella, and courts have long assessed them under a common proportionality and reasonableness framework.

Treating new generative AI platforms as variations on a well-established theme is both technically correct and a wise doctrinal choice. It avoids the disruption that a bespoke framework would unnecessarily invite and preserves the analytical coherence that has long been in place.

Parties may use search terms to prefilter hits for generative AI review.

The plaintiffs challenged LinkedIn's decision to filter out documents that did not hit on search strings before documents were sent to Relativity aiR for generative AI review. The plaintiffs did not, however, dispute the appropriateness of the search strings. Rather, they argued that the approach improperly reduced the universe of documents reviewed by the AI system and risked excluding responsive materials. The court disagreed.

Because the plaintiffs raised no challenge to the search strings themselves, the use of the strings to pre-cull the document population before AI-assisted review did not present a responsiveness problem. In so ruling, the court expressly relied on TAR cases holding that search terms may be used to reduce review populations before technology-assisted review is conducted.

This ruling reinforces a practical reality of modern discovery — using generative AI does not eliminate proportionality concerns. Parties may continue to employ reasonable filtering methodologies, including negotiated search terms, to reduce large data volumes before undertaking AI-assisted review.

The holding also aligns with the growing line of authority recognizing that search-string filtering before TAR is consistent with the reasonable inquiry standard of Rule 26(g) of the Federal Rules of Civil Procedure. Counsel who negotiate defensible search terms and layer TAR — generative or otherwise — atop the resulting document population act well within the certification obligations Rule 26(g) imposes on every signed discovery response.

Discovery on discovery remains disfavored.

When a party discloses that it plans to use a generative AI review tool, the opposing side often seeks detailed metrics such as error rates, elusion estimates and protocols for human validation review. Rejecting what the plaintiffs were demanding from LinkedIn, the court reiterated the familiar principle that discovery on discovery is disfavored, and appropriate only where a specific deficiency has been identified. Mere speculation that the process may be inadequate is not enough, even where the producing party plans to use generative AI.

This ruling offers useful authority for parties facing requests for AI-related metrics. To the extent parties want to ensure access to such information, they should consider including particularized disclosure requirements in ESI protocols before discovery begins. But demanding a stringent validation protocol for AI-assisted review often misses the foundational point that an appropriate discovery workflow functions as an integrated whole,

not in isolation.

Well-designed workflows embed quality control across the full life cycle precisely to eliminate the need for invasive, after-the-fact incursions into the producing party's methodology. Restraint around compelled disclosure of granular AI metrics — that is, requiring a demonstrated deficiency before allowing an audit into the AI methodology — is best understood as confidence in the workflow as a whole, rather than indifference to accuracy.

Conclusion

The decision in *Schulte v. LinkedIn Corp.* reinforces that existing discovery rules can and should accommodate emerging technologies. The case represents a meaningful step toward thoughtful adoption of generative AI in discovery, which should give litigants confidence that well-designed AI-enabled workflows can satisfy reasonableness and proportionality standards while delivering the efficiencies necessary in modern discovery.

Read together with other recent discovery-related authority, a broader theme emerges. Courts and the discovery community should converge on the shared premise that generative AI is not an exception to the discovery rules, but the next chapter in their evolving application.

Practitioners who invest in well-architected workflows, document their choices and negotiate protocols early will find the doctrine on their side. Those who treat AI as a novelty requiring custom-built rules will inevitably find themselves relitigating yesterday's questions.

Christopher Davies is a partner, Marissa Wenzel is counsel, and Zviad Guruli is director of ediscovery and data analytics at WilmerHale.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.