

Expert Opinion **Legal Technology**

TAR and Generative AI Team Up to Achieve Success in Large-Scale Discovery

The most successful discovery strategies often are built around the realities of a particular matter, not around technology for its own sake.

July 30, 2026 at 12:14 PM By **Robert Keeling** , **Amy Hanke**, **Allison Myers** and **Mark Resnick**



L-R: Mark Resnick & Allison Myers of Consilio, and Robert Keeling & Amy Hanke of Redgrave Courtesy photos

The Challenge: Too Much To Review, Not Enough Time

Large-scale regulatory investigations increasingly require producing parties to manage more data, tighter deadlines, shifting scope, and increased risks. In cross-border matters, the challenges can be even greater. For example, jurisdiction-specific requests can determine whether a document is responsive (e.g. sales documents from one, specific country) and where overbroad productions may expose the company to unnecessary risks.

This was the challenge in a recent non-U.S. antitrust investigation involving tens of millions of documents, an eight-week production timeline, and complex regulator requirements. To meet

these demands, the company needed a solution that was fast, scalable, defensible, cost-effective, and precise.

In this case, the solution was not TAR or gen AI. It was TAR and gen AI, deployed in sequence. TAR supplied the speed and scale needed to triage millions of documents. Gen AI then served as a precision-enhancing filter for the subset of documents TAR identified as potentially responsive. The results were compelling: in a 3.6 million-document population, **gen AI eliminated 27.5% of documents it analyzed from the responsive TAR set**—more than one in four documents that otherwise would have moved toward production consideration.

The Solution: A Layered Review Strategy

As collection and processing progressed, the review population totaled more than 69.6 million documents. Targeted culling reduced the population to about 5.5 million documents and email threading brought it down to about 4.1 million documents. TAR was applied to about 3.6 million documents.

Given the eight-week timeline, employing TAR was the logical first step because it could rapidly and more cost-effectively score the entire review population. The foreign regulator also required six issue-specific TAR models, each tied to its own custodians, date ranges, and subject matter. Because those scopes overlapped, the same document could fall within multiple models and raise different responsiveness questions depending on the issue.

Additionally, the document population was not static. Rather, documents continued to flow into the population throughout the compressed production window.

In these circumstances, our preferred approach was to layer gen AI over TAR, such that gen AI analyzed only the subset of documents where its contextual reasoning would add value by improving precision and reducing overproduction risk.

TAR as the Foundation

In today's AI-focused world, it may be tempting to overlook TAR. But this matter showed that TAR remains highly effective, especially when teams need to analyze a massive and complex document population quickly, consistently, and defensibly.

At the same time, TAR has a known tradeoff: a document that has a responsiveness score above the TAR cutoff may be a false positive, i.e., it is either not responsive or otherwise out of scope. In our example, TAR was identifying as responsive documents that otherwise met production criteria except for the fact that they were outside the jurisdictional scope, such as concerning only U.S. activities and not relating to the foreign jurisdiction where the document requests were focused. Thus, we needed more than TAR alone.

Generative AI as a Precision Layer

Gen AI added what TAR could not: contextual analysis against the jurisdictional scope of the investigation. The strength of the team's discovery strategy was **layering gen AI on top of TAR to combine the strengths of both technologies for maximum benefit**. While TAR quickly identified the documents most likely to matter, gen AI delivered important contextual analysis and greater precision to refine the results. Accordingly, combining the use of TAR for scale and speed with gen AI for context and nuance created a workflow that moved quickly through massive volumes while preserving defensibility, increasing accuracy, and reducing risk.

The workflow was straightforward. TAR-eligible documents were scored across six issue-specific TAR models, model-specific control sets established cutoff scores, and documents above the cutoff were treated as TAR-positive for that issue. Gen AI then analyzed those TAR-positive documents using prompts aligned to the relevant scope. A document moved forward as responsive only if it cleared both thresholds: it scored above the TAR cutoff and gen AI classified it as responsive. Responsive document families then went through privilege review.

The Result: Reduced Volume, Increased Precision, and Lower Risk

Of the 3.6 million documents subjected to TAR, about 1.8 million scored above the cutoff for at least one TAR model. Within that TAR-positive population, about 1.4 million documents were capable of being analyzed by gen AI.

Gen AI classified 375,675 (27.5%) of those documents as not responsive despite their TAR-positive scores. Thus, more than one quarter of the AI-eligible TAR-positive population was removed from the responsive set before family expansion. That was the core practical value of the layered approach. Because responsiveness decisions can drive family-level treatment, every nonresponsive document removed before family expansion can have a multiplier effect, reducing downstream review burden, production volume, and unnecessary risk.

Of the 375,675 TAR-positive documents that gen AI classified as not responsive, 350,981 were excluded based on issue scope. Critically, however, another 24,694 documents were excluded after gen AI prompts assessed whether the documents related to the regulator's specific national jurisdiction. Gen AI's capacity for contextual analysis allows it to exclude substantively similar documents that may be legally out of scope based on factors such as jurisdiction, product parameters, or the identity or role of the participants.

Defensibility Considerations

A layered workflow does not eliminate the ability to demonstrate transparency and defensibility, but it may change how teams perform validation. Although TAR control sets were used to determine cutoff scores, once gen AI is added, those control set statistics alone are no longer sufficient to statistically validate the end-to-end process. If formal validation is required, the discovery team should design a validation method that accounts for the combined workflow, such as a stratified sampling approach drawing from documents treated as both responsive and nonresponsive.

Key Lessons Learned

- **Gen AI is a powerful addition to the discovery toolset.** Gen AI brings capabilities in addition to TAR, particularly where responsiveness turns on nuanced contextual analysis. When deployed thoughtfully, gen AI can materially improve precision without abandoning the scalability of TAR.
- **Gen AI deployment should match operational realities.** Although gen AI capabilities continue to advance, in matters involving massive populations, rolling data ingestion, and tight deadlines, TAR remains an important consideration.
- **Gen AI adds value with sophisticated contextual analysis.** Gen AI was effective at identifying documents that look similar but differ in legally significant ways. This level of nuance is particularly helpful in allowing more precise responsiveness determinations in global investigations that have a specific jurisdictional scope.
- **Layering technology makes proportionality more practical.** Applying gen AI to the documents TAR determined were most likely responsive reduced cost, improved precision, and helped manage risk.

Conclusion: Consider Combining TAR and Gen AI

This matter showed that in large-volume reviews, TAR plus gen AI can be an effective workflow. The most successful discovery strategies often are built around the realities of a particular matter, not around technology for its own sake.

This investigation showed what is possible when established TAR workflows and gen AI are layered together. TAR delivered the scale and speed to manage millions of documents and a continuously evolving data set. Gen AI delivered the contextual precision to reduce noise and avoid unnecessary production. Together, they produced a workflow that was practical, defensible, and more cost-effective than employing either approach alone.

Robert Keeling *is a co-managing partner of Redgrave and a nationally recognized authority on eDiscovery. He serves as discovery counsel in complex, data-intensive matters, including litigation, investigations, and regulatory reviews.*

Amy Hanke is counsel at Redgrave. Her practice focuses on eDiscovery in complex litigation, regulatory actions, and investigations, including ESI protocol negotiation, technology assisted review and predictive coding, sensitive document reviews, and privilege issues.

Allison Myers is vice president of data analytics at Consilio, where she leads a team of 45 lawyers and programmers delivering custom analytics solutions for large, complex, and data-intensive eDiscovery matters. A frequent CLE presenter with more than twenty years in eDiscovery, she combines a legal education with technical expertise to help clients and counsel turn document populations into actionable intelligence.

Mark Resnick is senior director of AI and analytics consulting at Consilio with more than 15 years of experience helping legal teams apply advanced analytics and AI to complex eDiscovery matters. He works with clients to design, validate, and defend technology-driven approaches to document review and investigations that improve efficiency and reduce cost, and he has played a key role in Consilio's adoption of generative AI and other emerging technologies.

Page printed from: <https://www.law.com/legaltechnews/2026/07/30/tar-and-generative-ai-team-up-to-achieve-success-in-large-scale-discovery/print/>

NOT FOR REPRINT

© 2026 ALM Global, LLC, All Rights Reserved. Request academic re-use from www.copyright.com. All other uses, submit a request to asset-and-logo-licensing@alm.com. For more information visit [Asset & Logo Licensing](#).