

CASE STUDY

Automating Privilege Log Drafting with Generative AI

How HaystackID® Delivered 230,000 Regulator-Ready
Privilege Log Entries in Eight Business Days

HAYSTACK®

Background Overview

This case study describes how outside counsel and HaystackID partnered to support a multinational corporation in a regulated industry during a complex regulatory inquiry requiring extensive document production. The objective was to produce a clear, specific, and defensible regulator-grade privilege log **within eight business days**, while productions proceeded in parallel. The matter was tied to a transaction in a highly regulated industry requiring rigorous regulatory review.

Outside counsel guided the overall strategy for the regulatory response, including engagement with regulators, scoping of custodians, and supervision of rolling productions, while **HaystackID operationalized the technical workflows required to meet aggressive deadlines**.

The team employed a **human-guided GenAI approach** that directed the AI to select from predefined options for consistency and supplemented this with generated controlled narrative text to provide contextual detail.

Data was collected from approximately **95 custodians**, with final collections for a narrowed list of approximately **80 custodians** conducted approximately three weeks prior to the production deadline. In the five days leading up to the final deadline, the team produced approximately **1.5 million documents** and approximately **15 million pages**, contributing to a total of approximately **5 million documents** produced overall.

From an enterprise-wide collection of approximately **15 million documents**, approximately **3.5 million** remained post-search for review. Within this set, approximately **230,000 items were flagged for privilege consideration** and required description drafting.

To meet this challenge, outside counsel and HaystackID collaborated on a human-guided, GenAI-enabled workflow built and managed by HaystackID's Legal Data Intelligence team. This workflow combined structured prompts, controlled narrative text, and attorney validation to produce a regulator-ready privilege log in just eight business days. **The regulator accepted the submission without follow-up inquiries on descriptive elements.** The practices are transferable to investigations and litigation matters facing similar volume and schedule pressures.

Executive Summary

HaystackID and outside counsel supported a multinational corporation responding to a regulatory inquiry that required expansive productions and a comprehensive privilege log. Using the **HaystackID Privilege Review and Log Automation offering powered by eDiscovery AI**, draft privilege log entries were produced in a hybrid format that pairs structured elements with controlled narrative text. Attorneys validated and refined each entry through a separate editable field with full audit export. The tool produced approximately **230,000 entries across eight business days**, reaching an output of approximately **100,000 GenAI-drafted entries a day** during peak periods. The regulator accepted the privilege log with no follow-up inquiries on descriptive elements. These outcomes were achieved while managing rolling data refreshes and late-stage custodial expansion.

Matter at a Glance

15M

Enterprise
Collection

3.5M

In Scope
for Review

230K

Flagged for
Privilege

8 Days

Production
Timeline

Regulatory Setting

Complex regulatory inquiries require broad collections, rigorous culling, and detailed privilege logs. Regulators expect precise descriptions that clearly articulate the basis for privilege, along with **operational consistency, traceability, and auditability**. Meeting these expectations at scale necessitates controls that withstand legal scrutiny and facilitate efficient review.

Matter Narrative Scale and Cadence

The matter began with an enterprise-wide collection of **~15,000,000 documents**. After application of search terms, approximately **3.5 million documents remained in scope for review**. Within this population, approximately **230,000 were flagged for privilege consideration** based on curated attorney and law-firm name lists and domain indicators. Operations proceeded under parallel deadlines, with privilege logging executed alongside productions. Cadence pressures included rolling data refreshes and late-stage custodial expansion, requiring a workflow capable of absorbing additional inputs without sacrificing consistency.

Method

Human-Guided AI with Attorney Validation

The workflow unfolded in **three phases**.



Phase 1—Identification

Reviewers identified documents for privilege consideration leveraging curated attorney and law-firm lists and domains. The flagged population of approximately **230,000 documents was batched for AI-assisted drafting**.



Phase 2—Drafting

HaystackID experts and outside counsel evaluated prompt results ranging from fully structured outputs to free-form narratives. The selected hybrid format leveraged structured pick-lists, where prompt instructions directed the AI to select from a **defined set of options for advice context and privilege basis**, followed by a controlled narrative segment that captured nuance while avoiding disclosure of privileged substance. Prompts and style rules were refined to align with client preferences.



Phase 3—Validation and Auditability

Within the Relativity platform, the AI-generated text was stored in a read-only field to preserve an immutable record of the draft. Attorneys conducted quality control in a separate editable field, applying sampling to confirm consistency and targeted overrides for risk categories. All reviewer actions and dispositions were **captured for audit export, supporting defensibility and process transparency**.

Results

Performance, Outcomes, and Reconciliation

Across **eight business days**, the team reviewed the approximately **230,000 documents identified** for privilege consideration, conducting human validation of privilege classification entries and review of AI-generated log entries. AI-generated log entries throughput peaked at approximately **100,000 auto-generated entries in a single day**.

Following attorney validation, approximately **118,000 documents were confirmed as privileged**, and approximately **120,000 were overturned as not privileged**. Confirmed plus overturned totaled approximately **238,000, with variance** reflecting de-duplication, family linking, and QC exclusions consistent with regulatory review practice.

Regulator response: The regulator accepted the privilege log with no follow-up inquiries regarding descriptive elements.

Observations

Narrative Quality, Error Patterns, and Controls

Narrative quality and consistency. The hybrid structure yielded entries that consistently identified legal actors, explained the context of legal advice, and stated the basis for privilege in controlled, fluent language. Structured components drove uniformity, while the narrative segment supplied the specificity expected by regulators.

Reviewer confidence and focus. Uniform draft entries reduced the variation common to manual drafting. Attorney time concentrated on role attribution and mixed-purpose communications, with edits typically addressing edge cases rather than systemic issues.

Error patterns. Two themes accounted for most edits:

1. **Attorney-list inaccuracies**, including shared surnames and similar addresses, prompting list corrections and disambiguation rules using domains and titles.
2. **Mixed business/legal content**, where interwoven communications required targeted checklists and second-level QC to maintain consistent privilege rationales.

Controls and auditability. The project employed prompt versioning and change logs, read-only storage of AI-generated text, separate editable fields for reviewer changes, sampling with targeted quality control review, and audit exports documenting reviewer actions and final dispositions. Processing occurred within Relativity, and client data was not reused for model training.

Post-Mortem Metrics on GenAI Privilege Classifications

After human review outcomes were finalized, a post-mortem review was conducted to assess how GenAI's privilege classifications would have compared if they had been used for privilege determinations. Throughout the review process, GenAI was used solely to generate log descriptions and was **not leveraged for privilege determinations**.

Following attorney validation, approximately **118,000 documents** were confirmed as privileged, and approximately **120,000 documents** were overturned as not privileged. Of the confirmed privileged documents, GenAI accurately classified approximately **86% of them**.

It is important to note the **family effect inherent in privilege determinations**. Because GenAI assessed each document individually, some family members within otherwise privileged sets appeared as not privileged on their face. In practice, however, these documents were withheld as part of privileged families.

Recommendations Practices for Repeatable Success

Integrate earlier in the cycle. Running description generation in parallel with privilege coding smooths workloads, accelerates feedback loops, and reduces end-of-cycle pressure, especially when data continues to roll in.

Adopt attachment-aware style logic. Emails and attachments often benefit from different treatments. A practical pattern is narrative-style entries for emails and structured pick-lists for standalone attachments, with automated detection rules applying the appropriate style by document type.

Formalize legal-actor precision. Because privilege often turns on who is acting in a legal capacity, early documentation of disambiguation rules for common names and dual-role in-house counsel improves accuracy. Domain and title heuristics, reinforced by curated lists, reduce repetitive edits.

Conclusion A Scalable, Defensible Framework for Privilege Logging

This regulatory matter shows that a **human-guided GenAI workflow**—anchored by structured prompts, controlled narrative text, and attorney-led validation—can meet regulator expectations for speed, clarity, and defensibility in privilege logging. Working within Relativity and leveraging HaystackID Privilege Review and Log Automation, the team produced approximately **230,000 entries in eight business days**. Acceptance by the regulator without descriptive follow-up highlights the value of audit-ready storage, reviewer overrides in separate fields, sampling discipline, and attachment-aware logic. Organizations facing similar demands can apply these practices to scale description-heavy work while maintaining consistency and control.

Learn More Today.

Learn how HaystackID can help your team implement scalable, defensible workflows for privilege review and AI-assisted logging. [Connect with us](#) to explore practical solutions for complex legal, compliance, regulatory, and cyber matters.

About HaystackID®

HaystackID solves complex data challenges related to legal, compliance, regulatory, and cyber requirements. Core offerings include Global Advisory, Cybersecurity, Core Intelligence AI™, and ReviewRight® Global Managed Review, supported by its unified CoreFlex™ service interface. Recognized globally by industry leaders, including Chambers, Gartner, IDC, and Legaltech News, HaystackID helps corporations and legal practices manage data gravity, where information demands action, and workflow gravity, where critical requirements demand coordinated expertise, delivering innovative solutions with a continual focus on security, privacy, and integrity. Learn more at HaystackID.com.

Assisted by GAI and LLM technologies.