

Built for Corporate Restructuring: The Case for Purpose-Built AI

CONTENTS

About Stretto

SECTION 1 Five Realities Reshaping How
Corporate Restructuring Gets Done

SECTION 2 Why Restructuring Demands
Purpose-Built AI

SECTION 3 From Concept to Practice: AI at
Work in Restructuring

SECTION 4 Deploying AI Responsibly in
High-Stakes Environments

SECTION 5 What Comes Next: The Evolution of
AI in Restructuring



About Stretto


Stretto delivers a full spectrum of case management and claims administration services, depository and distribution solutions, and technology tools to legal and financial professionals. With a comprehensive suite of tailored offerings, Stretto provides an unparalleled portfolio designed to meet our clients' unique financial and business objectives. For more information about Stretto, please visit stretto.com

SECTION — 1

Five Realities Reshaping How Corporate Restructuring Gets Done

Corporate restructuring operates at the intersection of legal precision and financial complexity, with stakeholder decisions routinely made under exigent circumstances. More data and shortened timelines have made that work harder. A 2025 study by Harvard and USC law professors published in the University of Chicago Law Review found that the average duration of large business bankruptcies declined dramatically over the past four decades, from approximately three years to three months. The attorneys, financial advisors, and trustees navigating these matters are asked to deliver faster, more defensible outcomes with the same or fewer resources. AI is the most practical tool available to meet that pressure.

Average duration of large business bankruptcies from filing to confirmation (in days)



Source: Roe & Simkovic, *Bankruptcy's Turn to Market Value*, 92 U. Chi. L. Rev. 285 (2025).

SECTION — 1

For purposes of this discussion, “AI” refers primarily to generative AI and related large-language-model tools, not to rules-based automation or traditional predictive models. The relevant use case in restructuring is the ability to analyze, summarize, compare, extract, and draft from large volumes of structured and unstructured information.

Not all AI performs equally in high-stakes environments, and the differences matter.

Five realities define where AI stands in restructuring today.

- 01** AI has moved past the experimentation stage. Firms and administrators that treated AI as a curiosity two years ago are deploying it operationally today. The share of legal organizations actively integrating AI nearly doubled between 2024 and 2025, according to Thomson Reuters, and that trajectory is continuing.¹
- 02** The problem is not simply that broad-market AI tools are generic. It is that restructuring work depends on contextual understanding across expansive, messy datasets, while the output must remain precise, stable, and defensible. The process is probabilistic, but the required output is much closer to deterministic. That is where off-the-shelf tools tend to break down.
- 03** Purpose-built AI designed around restructuring-specific workflows, documents, and language improves both precision and traceability. A system grounded in thousands of DIP orders, critical vendor motions, and plan confirmation hearings recognizes the language and structure of that work in ways a general-purpose model cannot. Precision means outputs grounded in the actual record of the case and comparable precedents, not generalizations drawn from broadly similar legal content. Traceability improves because the system is designed to recognize which documents are authoritative, account for how filings supersede one another, and cite every output to a specific source for practitioner review.

¹ Thomson Reuters Institute, *Future of Professionals Report*, 2025.

SECTION — 1

- 04** Security, confidentiality, and auditability must be foundational. Generative AI introduces risks that traditional legal technology did not. Inputs submitted to general-purpose models may be retained, used for training, or accessible outside the session. For restructuring professionals working with privileged communications, unreleased financial data, or personally identifiable information, that is a material concern. Auditability requires the ability to demonstrate what information the AI system was given, what it returned, and which model and session produced it. In court-supervised proceedings, accountability is a professional requirement.



The attorneys, trustees, and organizations that can synthesize complex information faster and act on it with confidence will define the standard.



- 05** The next competitive advantage in restructuring will come from time-to-insight. Thomson Reuters projects that AI-enabled legal professionals stand to reclaim roughly 240 hours annually.² That time translates directly into faster analysis and more defensible outcomes. The attorneys, trustees, and organizations that can synthesize complex information faster and act on it with confidence will define the standard.

² Thomson Reuters Institute, *Future of Professionals Report*, 2025.

SECTION — 2

Why Restructuring Demands Purpose-Built AI

Restructuring operates differently from any other area of legal or financial practice. It straddles litigation and transactional work, and unlike traditional M&A or private financings that happen largely behind closed doors, restructuring plays out in public through the court system, on the record, with every motion, objection, and order available for scrutiny. The average large, complex Chapter 11 case will generate more than 1,000 docket entries just in the first 12 months of the case, with the most active cases exceeding 3,000 docket entries in the first year. Over the full lifespan of a complex case, docket sheets routinely exceed 10,000 entries.

10,000+

docket entries over a complex case's full lifespan

1,000 docket entries
avg. case, first 12 months

3,000+ docket entries
most active cases, first 12 months

Yet most of that data is effectively walled off. While PACER provides some public access, the bulk of its documents and search results sit behind a paywall, and content on claims agent platforms is difficult for general-purpose AI tools to access. Creditor data, claims information, and case administration records present the same problem. This leaves gaps in the training data of any tool not built specifically for this work. For generative AI systems, which derive their usefulness from pattern recognition across large training datasets, those gaps are consequential. General-purpose models trained on broad legal and financial content rather than restructuring-specific filings, motions, and orders produce confident-sounding output where generic concepts and inapplicable legal frameworks parade as domain expertise. The Lehman Brothers docket has nearly 63,000 entries and only 24 published opinions. The Purdue Pharma docket tells the same story, with over 9,000 entries and

SECTION — 2

only six published opinions. The high-value, day-to-day work of restructuring lives in the motions, responses, orders, and stipulations that never make it into a reported decision. General-purpose AI simply cannot see most of that record.

Purpose-built AI starts with a dataset built around restructuring-specific material and removes what is irrelevant. Equally important is how the system is designed to work with that data. Restructuring cases generate documents in sequence, and the first version filed is rarely the operative one. A system that cannot distinguish a proposed order from the final entered version, or an interim DIP order from the amended and restated final one, will surface the wrong answer with the same confidence as the right one. Domain-specific architecture accounts for that document hierarchy, understanding not just what was filed but its sequence and whether it was superseded. With that architecture in place, practitioners can confidently synthesize staggering volumes of case data, from DIP financing negotiations to valuation disputes to Section 363 sale processes, across cases and in a fraction of the time previously required.



When professionals can see more of the record, within a case, across comparable cases, and across a judge's full body of work, they make better-informed arguments.



Representing a party in a retail restructuring? Purpose-built AI can surface how a plan in a comparable case evolved from initial proposal to final confirmed version: what critical vendor relief looked like, how stalking horse bid protections were structured, and what DIP terms the market accepted. It can highlight insights into a judge's actual Chapter 11 track record beyond published opinions, including rulings across contested and uncontested matters in cases where that data is available.

The domain specificity requirement does not change based on the task. The same logic applies across every dimension of restructuring practice, from creditor communication to document security to case research. When professionals can see more of the record, within a case, across comparable cases, and across a judge's full body of work, they make better-informed arguments.

SECTION — 3

From Concept to Practice: AI at Work in Restructuring

The gap between what AI promises and what it delivers comes down to whether it was built for the work. Stretto Intelligence is Stretto's innovation engine, developing and applying AI where it delivers the greatest value across the complex, high-stakes work of restructuring and bankruptcy professionals. Every tool, insight, and service is purpose-built for the day-to-day realities of insolvency proceedings and meets the highest standards of security and confidentiality. The three applications below illustrate what that looks like in practice.

Research Suite by STRETTO

Every new matter starts with the same question: where do you begin? For a junior associate staffed on a Chapter 11 debtor representation for an airline based in Florida, the answer is precedent. But precedent research in restructuring is rarely straightforward. You know the cases your firm has handled and the ones that made headlines. Neither helps when your firm has never represented an airline, has never filed a case in Florida, and the new matter has no obvious analog in your experience. Identifying which filings are relevant to a specific matter, in a specific jurisdiction, with specific operational and capital structure characteristics, is itself a substantial undertaking, and that is before you have reviewed a single document.

Research Suite by Stretto addresses that problem. Its database covers a curated collection of more than 4,000 of the largest Chapter 11 cases and over 5.5 million court filings, with approximately 200 significant new cases added each year, all searchable through filters built specifically for restructuring workflows, not adapted from general litigation tools. From a broad universe of potentially relevant cases, practitioners can quickly identify the most comparable precedents, then use proprietary bankruptcy-specific metadata filtering to zero in on the specific filings that matter. The filters surface every critical vendor motion filed in comparable airline bankruptcies and every DIP financing order entered in the Southern

SECTION — 3

District of Florida over the last five years. The Precedent Package's AI Dossier then transforms that curated set into a comparative memorandum with an executive summary, case-by-case takeaways, side-by-side analysis of facts, legal arguments, and outcomes, with every sentence cited to the specific page of the underlying filing.

Work that once took hours or days now produces a verifiable, fully-cited first draft in minutes. In one recent instance, Research Suite generated a complete Precedent Package for a collection of 80 pleadings relating to first day vendor relief from ten large airline cases in 25 minutes, reviewing 985 pages of court filings and producing 80 individual document summaries, ten case summaries, and a 26-page memorandum-style AI Dossier.

4,000+

of the largest
Chapter 11 cases

5.5M+

court filings

~200

significant new
cases added yearly

Source: Research Suite curated database

With Research Suite, practitioners spend less time locating and assembling information, and more time doing the work that drives outcomes: developing strategy, sharpening arguments, and crafting more persuasive pleadings. Better information flowing into the strategic portion of the work makes every hour spent there more productive.



In large Chapter 11 cases, the administrative burden of creditor communication is substantial and often underestimated. Major consumer-facing bankruptcies often involve creditor populations ranging from 50,000 to hundreds of thousands of individuals, many with no legal background, each needing timely access to accurate case information: key deadlines, plan developments, what a particular filing means for their claim. Celsius had more than 1.7 million registered users and approximately 300,000 active users with substantial outstanding balances at the time of filing.³ Similarly, Azul, one of Brazil's largest airlines, faced approximately 85,000 litigation claims when it filed for bankruptcy, in addition to thousands of other non-litigation creditors.⁴

³ Stretto CORE, proprietary claims-management platform for Celsius Network LLC, U.S. Bankruptcy Court, Southern District of New York, Case No. 22-10964, filed July 13, 2022.

⁴ Stretto CORE, proprietary claims-management platform for Azul S.A., U.S. Bankruptcy Court, Southern District of New York, Case No. 25-11176, filed May 28, 2025.

SECTION — 3

Historically, meeting the communication needs of creditor populations at that scale required significant staffing resources, and responses were rarely immediate. Creditors navigated dense legal filings on their own, called into support centers, and waited. In cases with large creditor populations or heightened public interest, that friction compounds quickly, creating bottlenecks that strain case teams and leave stakeholders without the clarity they need to make informed decisions.

Stretto Conductor solves that problem by embedding an intelligent, always-available assistant directly into case websites. Creditors can ask questions in plain language and receive clear, tailored answers in real time, without navigating technical court documents or waiting on traditional support channels. For filings that would otherwise require legal expertise to interpret, Conductor generates concise summaries that distill key points, deadlines, and implications into language accessible to a non-specialist audience, with integrated citations to the relevant pleadings.



Conductor generates concise summaries that distill key points, deadlines, and implications into language accessible to a non-specialist audience, with integrated citations to the relevant pleadings.



Accurate responses depend on Conductor understanding how filings relate to one another, including conflicting provisions, multiple proposed orders, and interim versus final rulings. A creditor asking about the status of a bar date needs an answer grounded in the operative court order, not an earlier proposed version, a superseded interim ruling, or language from an exhibit attached to a prior motion. Conductor is also designed to recognize when a question crosses from factual inquiry into legal advice and respond accordingly, maintaining appropriate boundaries without disrupting the user experience. Rather than adapting a generic chatbot, Stretto built Conductor from first principles and has been granted a patent for the solution. The platform scales to meet the demands of

SECTION — 3

each case without requiring proportional increases in staffing, support resources, or costs to the estate. Conductor is now deployed across over 100 case websites and has facilitated more than 100,000 conversations since its launch in 2025, handling thousands of interactions each month.

The impact is clearest in cases where the volume of creditor activity is highest. In the Saks Global Enterprises retail restructuring, Conductor handled over 4,000 chat sessions and delivered more than 3,000 docket summaries in just the first two months of the case. Case teams fielded fewer repetitive inquiries, and creditors received faster, more consistent access to the information they needed. That combination of scale, speed, and accessibility is changing how restructuring administration works.



Restructuring professionals handle some of the most sensitive financial and legal information in practice. This data carries real confidentiality obligations and fiduciary risk. Yet many of the AI tools available today introduce a risk that practitioners cannot ignore: providers may retain inputs, share them, or use them to train the underlying models. It is telling that in the 2026 Wolters Kluwer Future Ready Lawyer Survey, data privacy and ethical concerns ranked as the single largest barrier to AI adoption among legal professionals, cited by 39% of respondents, ahead of training gaps, integration challenges, and resistance to change.⁵ For Chapter 7 trustees and restructuring teams working with personally identifiable information and confidential client materials, that risk is real and immediate. It is a meaningful barrier to using AI at all, and it forces a choice between the efficiency of the technology and the fiduciary obligations of the work.

Stretto SecureAI removes that barrier by providing an encrypted, zero-retention platform where practitioners can use advanced AI models without concern that their inputs will be stored, shared, or used for any purpose beyond the immediate session. Each user operates within a private workspace, protected by individually scoped access controls and enterprise-grade encryption. Decryption permissions are controlled at the user level, meaning access to any given workspace depends on the explicit authorization of the user who owns it, not simply platform-level credentials. Those retention controls extend to the third-party model providers integrated into the platform, closing the gap that most general-purpose tools leave open when user data passes through external model APIs. A built-in prompt library gives users a growing collection of restructuring-specific prompts to draw from, helping practitioners of any experience level quickly get results. Professionals can upload petitions,

⁵ Wolters Kluwer, *Future Ready Lawyer Survey*, 2026.

SECTION — 3

schedules, and case filings and generate tailored summaries calibrated to the specific details most relevant to their responsibilities, such as exemptions, gaps in petition data, and creditor patterns. Legal teams can conduct research and analysis that would otherwise require relying on tools not designed for the confidentiality standards legal practice demands.

When data exposure is no longer a concern, the technology can be applied to the full scope of the work rather than only the portions deemed safe enough to share. In just the first six months, SecureAI users completed thousands of sessions, accounting for over 100 million tokens of LLM usage. As with all AI-assisted work, SecureAI's terms make clear that users are responsible for reviewing and evaluating outputs before relying on them for advice or case decisions, a standard that reflects the professional obligations the platform was built to support. Legal teams focus on analysis and judgment rather than compliance anxiety. The efficiency gains that AI promises in restructuring become accessible without compromising fiduciary and privacy standards.



Legal teams can conduct research and analysis that would otherwise require relying on tools not designed for the confidentiality standards legal practice demands.



SECTION — 4

Deploying AI Responsibly in High-Stakes Environments

AI tools used in corporate restructuring must operate inside a framework that protects confidentiality, supports review, and holds up under scrutiny. Handling the financial, legal, and stakeholder information that restructuring work involves demands clear limits on how sensitive data is used, encryption in transit and at rest, permission-aware access controls, and retention policies. According to IBM's 2025 Cost of a Data Breach Report, U.S. organizations faced an average breach cost of \$10.22 million, the highest of any country studied, driven by steeper regulatory fines and increasingly complex breach investigations.⁶

Outputs from these tools must be reviewable and traceable. Responses grounded in approved sources and supported by citations give professionals the ability to verify the basis for an answer before it informs client advice, case strategy, or court-facing work. The consequences of failing to verify AI outputs are well documented and escalating. In *Mata v. Avianca*, attorneys were sanctioned for submitting fabricated AI-generated citations to a federal court, which was the first major incident to trigger widespread judicial and regulatory attention. More recently, in April 2026, restructuring attorneys at an AmLaw 100 firm filed an emergency motion before the U.S. Bankruptcy Court for the Southern District of New York containing AI-generated errors including fabricated citations and misquoted authorities. The firm filed a letter apologizing to the court.

These are enforceable professional obligations. In July 2024, the ABA issued Formal Opinion 512, its first formal ethics guidance on generative AI, confirming that existing Model Rules apply directly to AI use in legal practice. Model Rule 1.1 requires lawyers to understand the capabilities and risks of the technologies they deploy. Model Rule 1.6 obligates them to protect client confidentiality regardless of what tool is used to handle that information. Model Rule 5.3 extends supervisory responsibilities to the use of AI tools, just as it does to any nonlawyer assistant's work. For restructuring practitioners, source citations, zero-retention architecture, audit logging, and tenant-scoped access controls are the mechanisms through which responsible AI deployment satisfies these obligations.

⁶ IBM Security, *Cost of a Data Breach Report, 2025*.

SECTION — 4

Responsible purpose-built AI reflects design decisions that general-purpose tools never confront. General-purpose tools generate responses by drawing on a broad training corpus, a vast largely unstructured body of text scraped from across the internet and legal databases, with no guarantee that restructuring-specific content was meaningfully represented. Purpose-built systems retrieve from a defined, curated dataset of restructuring filings, orders, and precedents. That constraint makes source-level citation possible: the system knows what it is drawing from because the dataset was deliberately assembled and bounded. Model governance adds another layer of accountability. A model update can change how a system interprets documents, structures outputs, or handles unfamiliar fact patterns, often without any visible signal to the user. Disciplined deployment means every model version in use is documented, that changes are evaluated against known benchmarks before rollout, and that the audit record captures which version produced any given output. Controlled rollout through feature flags allows new model versions to be introduced incrementally and rolled back if behavior deviates from expectations. Prompt governance warrants the same rigor, since prompts change more frequently than models and introduce their own variability. For practitioners operating under court supervision, that level of traceability makes AI use professionally defensible.



Purpose-built systems retrieve from a defined, curated dataset of restructuring filings, orders, and precedents.



AI can accelerate the work, but the practitioner owns the professional responsibility. AI deployment in legal practice requires human-in-the-loop review at every material decision point. The practitioner remains the final authority on work that will ultimately be tested in negotiations or in court.

What Comes Next: The Evolution of AI in Restructuring

The evolution already underway in restructuring practice is driven by AI that is increasingly capable of handling the specific demands of this work. Structured datasets, traceable outputs, and governed model behavior are the foundations that make the next stages of that evolution possible. Those stages will move across three distinct phases, and the industry is already in the middle of the first.

PHASE ONE: INTELLIGENCE-ENABLED ANALYSIS

The volume of court filings and other case data available in any major restructuring already exceeds what any team can synthesize manually in a compressed timeline. AI designed for restructuring makes that synthesis possible, surfacing relevant precedent, flagging structural comparisons, and delivering insight at the speed decisions need to be made. The result is the ability to do analysis of a depth and breadth that was not previously possible.

PHASE TWO: STRATEGIC CO-PILOT SYSTEMS

The tools practitioners use today for drafting and filing will evolve into systems capable of something more consequential. A strategic co-pilot system assists practitioners in modeling how a negotiation or proceeding is likely to unfold based on how comparable situations have resolved. This functionality enables them to stress-test positions, anticipate objections, and refine strategy before reaching the negotiating table or the courtroom. Outcome modeling is the core capability that makes this possible. It depends on training a model to recognize which conditions, across jurisdiction, fact pattern, creditor composition, and market environment, predict which results. Achieving that pattern recognition at the scale and specificity restructuring work requires is the foundational challenge that separates phase two from phase one.

SECTION — 5

PHASE THREE: INTEGRATED PLATFORMS

Right now, most AI tools operate in isolation. The next generation of restructuring intelligence platforms will connect claims management, legal research, drafting, and stakeholder analysis into unified systems where insight compounds across the full lifecycle of a case. Genuine integration requires shared data architecture across functions, so that research context informs drafting, claims data surfaces in negotiation analysis, and creditor communication patterns are visible to the teams developing case strategy. This can only be accomplished through architecture designed from the ground up to connect functions rather than aggregate tools. Each layer informs the next, and the resulting cumulative intelligence becomes something no single tool operating alone could produce.

The progression from intelligence-enabled analysis to strategic co-pilot systems to integrated platforms is not a straight line, and it will not happen uniformly across the industry. The firms and practitioners that engage with these purpose-built tools today will position themselves to adopt the next generation of systems as they emerge, and ultimately define what competitive and strategic restructuring practice looks like for the next decade.



About Stretto Intelligence

Stretto Intelligence is Stretto's innovation engine, developing and applying AI where it delivers the greatest value across the complex, high-stakes work of restructuring and bankruptcy professionals. Every tool, insight, and service in the portfolio is purpose-built for the day-to-day realities of insolvency proceedings that general AI was never designed to understand, and meets the highest standards of security and confidentiality. A dedicated team helps firms implement, adopt, and get real value from this technology, reflecting Stretto's broader commitment to staying ahead by deploying emerging technology with the precision and expertise this industry demands.

To learn more about Stretto Intelligence, please reach out to us at aiinfo@stretto.com

*This document is proprietary to Stretto.
The perspectives and frameworks
presented here reflect Stretto's
expertise and point of view on the
application of artificial intelligence in
corporate restructuring.*

