

Your Contracts Can Now Talk. Is Anyone Listening?

Law Department Management



Banner artwork by Butusova Elena / *Shutterstock.com*

Ask a contract repository a question like, "What is the aggregate liability exposure across every customer agreement signed since 2023?" and until recently the honest answer was to wait several weeks while outside counsel reviewed the portfolio by hand. That timeline is now collapsing.

Open protocols such as the Model Context Protocol (MCP), which let AI agents connect directly to governed document repositories, are turning that lengthy procedure into something closer to a conversation, and legal teams that have made this connection can run this kind of analysis between meetings.

This is worth applauding: our contracts can now “talk” and serve up answers to business-critical questions. So, what’s the problem?

The flip side is what happens when that same AI capability is pointed at a repository that is not ready for it. If contracts are inconsistently classified, metadata varies with each filer's judgment, and access permissions were never designed with this kind of use case in mind, connecting an AI agent industrializes governance debt.

That is the challenge now facing legal professionals. Organizations that spent the past several years

on classification, metadata discipline, and access governance are positioned to collect a real dividend from this new layer of connectivity. Organizations that have not are about to find out, fast, exactly how disorganized their contract data really is.

> Organizations that spent the past several years on classification, metadata discipline, and access governance are positioned to collect a real dividend from this new layer of connectivity.

How governance debt creeps in

Most contract repositories accumulate the same problems over time, largely invisibly, because each individual filing decision seemed reasonable in isolation.

Consider a single customer relationship. Its agreements might be filed under a formal legal entity name, a trade name, an abbreviation, a pre-acquisition name, and a foreign subsidiary: five labels for one commercial relationship. A human reviewer recognizes the connection without much effort. An AI system treats each as a distinct party, severing the link between a master agreement and its amendments and making it impossible to reason about the full scope of a relationship.

Or consider how several versions of the same document end up marked “final.” After a negotiation reaches an agreed form, one side requests a last-minute change, a revised version is saved under the same label, and the executed copy is uploaded without displacing either predecessor. An AI tool has no way to know which of the three is actually valid and will surface all of them with equal authority.

Fragmentation compounds both problems. Even legal departments that use a document management system rigorously often have contracts scattered across multiple tools and locations.

Access and permissions further complicate things: permission structures built for humans navigating a folder tree were not designed for an AI agent that can surface anything it can technically reach in a single query, so where permissions were set as an afterthought, a conversational interface only makes the consequences visible to more people, faster.



The proper sequence is IA before AI

While it might sound like telling people to eat their vegetables, the solution here is to shore up the information architecture (“IA”) before layering on AI. In other words: think “IA before AI.” That’s the proper sequence of events, and the human capital cost involved in organizing your existing data is now as important as the tool itself. As a planning rule of thumb, budget about five minutes of human review per agreement. For a 20,000-contract legacy corpus, that is roughly 1,700 hours, close to a full year of one person’s time.

No AI initiative functions reliably without a foundation, and building that foundation breaks down into four areas of work. None is technically complex, but all four require sustained organizational commitment, which is why they tend to get deferred in favor of more visible AI initiatives.

> No AI initiative functions reliably without a foundation, and building that foundation breaks down into four areas of work.

The first focus area is consolidation and data normalization: designating one governed repository as the source of truth for executed contracts. Then establishing and enforcing consistent standards for how contract data is entered, stored, and maintained. Ideally by embedding metadata requirements directly into the upload workflow, so that, for example, a system will not accept a new document without a counterparty name in the required format.

The second is taxonomy development: building a classification system that tells a repository, and any AI tool reading it, what kind of document it is looking at. An AI-ready taxonomy treats document status (draft, executed, expired, superseded) as a separate dimension from document type, rather than encoding status into a file name, and forces resolution of long-standing ambiguities, such as the practical difference between a master services agreement, a professional services agreement, and a consulting agreement, which can carry genuinely different obligations.

The third, and most frequently avoided, focus area is historical remediation: applying current standards retroactively to the legacy corpus rather than only to documents created going forward. Launching a clean intake process solves the problem for new documents but does nothing for the years of documents that already exist, and that gap does not shrink on its own.

The fourth is access governance: re-scoping permissions around what an agent can reach in a single query. That means least-privilege access, so an agent surfaces only what the requesting user is entitled to see; ethical walls and matter-level restrictions enforced at the repository layer; and audit trails that record what each agent retrieved, and for whom.

Progress over perfection

The instinct, when a legacy corpus feels overwhelming, is to defer this work. Don’t. Favor progress over perfection.

The way to make remediation manageable is to tier it by risk rather than attempting everything at once. Prioritize contracts above a material dollar threshold, active commercial relationships, agreements with automatic renewal provisions, or contracts governed by nonstandard or higher-exposure jurisdictions.

An automated enrichment run across a legacy corpus, even without a full human validation pass on every record, generally produces a baseline that is a marked improvement over searching manually across disconnected systems, provided human oversight remains in place before that data informs a decision.



Unlocking the value

With that work completed, organizations can begin extracting meaningful insights from their contract data. One area attracting significant attention is trend analysis across large volumes of agreements.

One area attracting significant attention is trend analysis across large volumes of agreements.

For example, AI can analyze changes made during negotiation and surface meaningful trends. Imagine that a company's standard template specifies Delaware as the governing law, yet in practice 99 out of 100 contracts are revised to use New York law instead.

Historically, discovering that pattern would have required a dedicated review project. The value comes when the system proactively surfaces the trend, freeing the legal team to focus on the business question: Is it worth slowing down negotiations over a provision that is changed almost every time? If not, the organization can update its template and reduce friction in future negotiations.

Another major area of potential value is obligations management. This goes beyond simple metadata extraction to focus on data tied to a required action, from tracking a renewal date to meeting compliance, security, or operational commitments.

For instance, a compliance team may want visibility into customers that negotiated unique security terms that could affect future changes to the company's security architecture.

Where to start?

Begin with a readiness test. Can you answer these three questions about your repository today, without launching a manual project?

1. How many active agreements, including amendments, do you have with your top 10 counterparties?

2. For any given contract, which version is the executed, governing one?

3. Who, and which AI tools, can see each agreement?

If any answer requires a manual review, start small with a 90-day pilot built around one question the business needs answered, such as the aggregate liability exposure question that opened this article. That question defines which contracts matter and what “organized” must mean for them.

Days 1 to 30: Scope and clear. Choose one risk tier that serves the question, such as active agreements above your materiality threshold. Map where those agreements live across every system, then remove duplicates and superseded drafts, checking retention schedules and legal holds before anything is deleted.

Days 31 to 60: Structure and secure. Adopt one file naming standard and apply it to the tier, resolve each counterparty to a single name, and add one mandatory status field (draft, executed, expired, superseded) enforced at upload. Then, review which users and systems can access these documents.

Days 61 to 90: Connect and measure. Connect an AI agent to the remediated tier and ask your anchor question. Audit answers against the contracts themselves. If the agent answers correctly in minutes, when the same question once required a manual project, you have your business case for extending the work to the next tier.

With intention and incremental work, your buried documents will soon become active business assets.

Join ACC for more AI insights!

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors’ employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

Cat Carroll



Head of Legal Operations

iManage.

Cat Carroll is Head of Legal Operations at iManage.