SUMATOSOFT

Industries: Legal services
Region: USA

PROJECT TYPE
Multi-agent contract review platform

TECHNOLOGIES
Python, LangGraph, FastAPI, PostgreSQL, Pinecone, Azure OpenAI (GPT-4.1), Anthropic Claude 3.5 Sonnet, React, TypeScript, Docker, Azure

DURATION
6 months

METHODOLOGY
Scrum

TEAM
1 AI Architect
2 ML Engineers
1 Backend Engineer
1 Frontend Engineer
1 QA Engineer
1 Business Analyst

sumatosoft.com
info@sumatosoft.com

Contract review multi-agent platform for a corporate law firm

A multi-agent contract review platform for a US law firm focused on tech-sector M&A and commercial contracting. The system extracts clauses, scores risk against firm precedent, and proposes redlines, with attorney review at every step.

Project background

The Client is a US-based mid-market law firm with 80 attorneys, focused on tech-sector M&A and commercial contracting. The firm runs a high volume of NDAs, MSAs, and SOWs across multiple practice groups, and most contract review is carried out by associates working from individual checklists and templates.

As deal volume grew, review consistency across associates became uneven, and the firm recorded a near-miss on an indemnity carve-out that almost derailed a transaction. Leadership decided to evaluate AI-assisted contract review with a structured assessment phase before committing to a build.

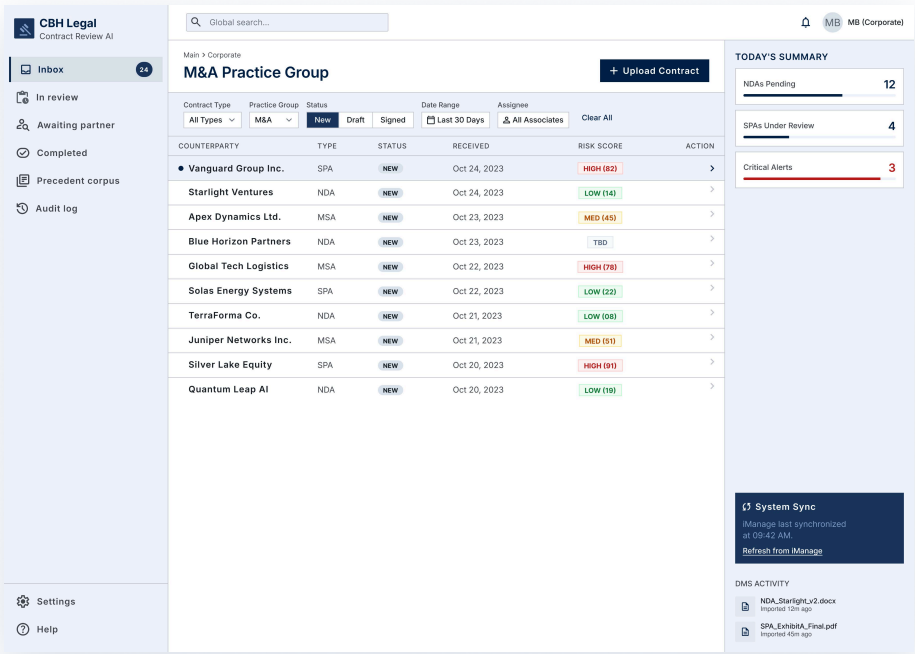

Project Distinctive Features

- ✓ Multi-agent architecture with separation of clause extraction, risk scoring, and redline drafting
- ✓ Precedent-grounded risk scoring against the firm's own contract corpus

SUMATOSOFT

“

SumatoSoft helped us put real structure around AI in our contract practice. The multi-agent design matched how our attorneys already work, and the assessment phase gave our partners a clear basis to commit. We have meaningfully cut review time without giving up the attorney oversight that defines our work.

CEO

- Attorney-in-the-loop review at every agent step
- Clause-level redline suggestions with rationale and citation back to source clauses
- Assessment phase preceding MVP development with documented evaluation criteria
- Support for NDAs, MSAs, SOWs, IP assignment, and employment contracts
- Full audit trail capturing every agent decision and attorney action
- Deployment inside the Client's Azure tenant to keep contract content within the firm's data perimeter

Business challenge

The Client wanted to reduce associate time spent on routine contract review and bring consistency to clause-level risk assessment without losing attorney oversight of every redline.

Additional challenges:

- Preserve attorney sign-off at every step of the workflow
- Ground risk scoring in the firm's own contract precedent rather than generic legal norms

Our solution

SumatoSoft delivered a multi-agent contract review platform that helps associates extract, score, and redline contract clauses against the firm's own precedent. The platform coordinates three agents: a clause extractor, a risk scorer, and a redline drafter, with attorney review checkpoints between each stage. Output flows into the Client's document management system as suggested redlines, never auto-applied changes.

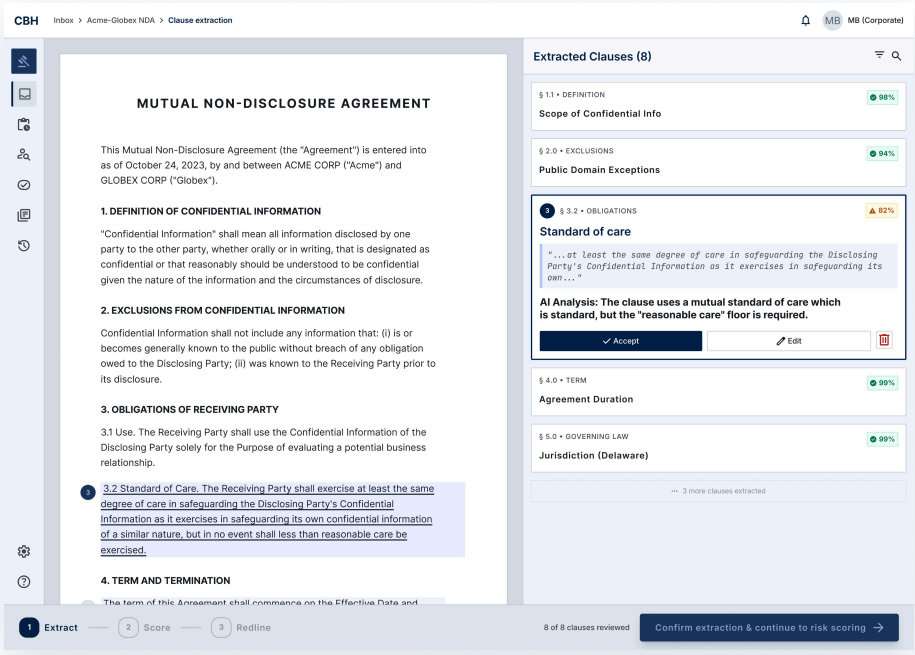

sumatosoft.com

info@sumatosoft.com

SUMATOSOFT

Assessment phase

Before any code was written, our business analyst and AI architect ran a four-week assessment with the Client's M&A and commercial contracting groups. We mapped the existing review workflow, identified the five contract types responsible for the majority of associate review hours (NDAs, MSAs, SOWs, IP assignment, employment), and audited a representative sample of 200 historical contracts together with the partners' redline notes. The assessment produced a documented use-case shortlist, an evaluation framework with attorney-validated success criteria, and an architecture proposal that the firm's executive committee approved before MVP development began.

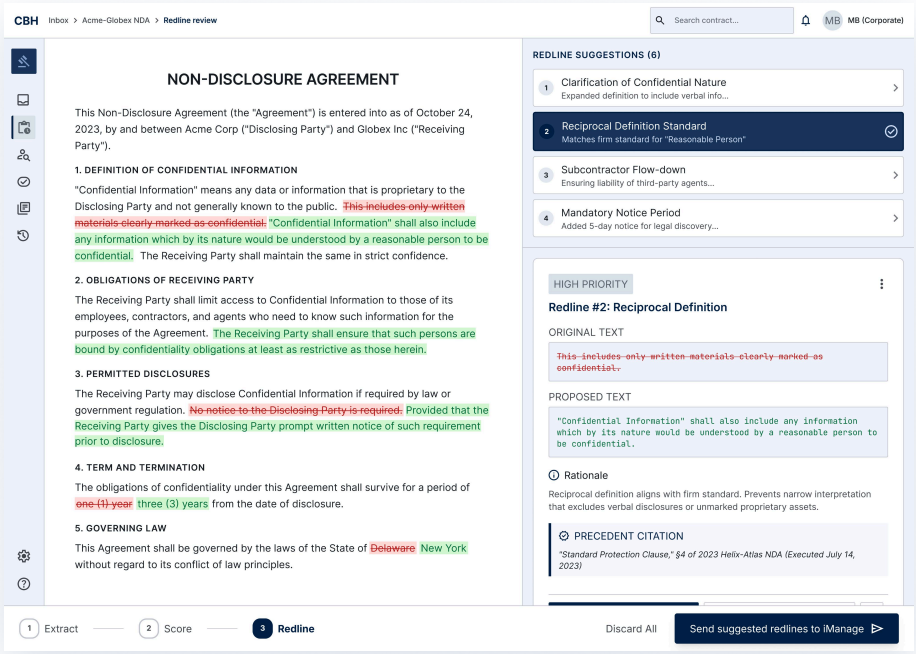

Multi-agent architecture

The three agents mirror the sequence associates already follow: read first, score second, draft third. We chose this over a single end-to-end model call so that attorney intervention is a first-class step in the graph (not an afterthought) and any agent decision can be inspected, overridden, or rerun without reprocessing the whole document. We built the orchestration layer on LangGraph. Clause extraction uses Azure OpenAI GPT-4.1 against a firm-tuned clause taxonomy. Risk scoring runs against the firm's precedent corpus indexed in Pinecone, retrieving comparable clauses and the partners' historical positions on each. Redline drafting uses Anthropic Claude 3.5 Sonnet, prompted with the firm's preferred language patterns extracted during the assessment phase.

Attorney-in-the-loop workflow

Every agent step produces a reviewable artifact in the web interface. Associates accept, edit, or reject each extracted clause before it moves to scoring; they can adjust risk levels with a justification before redlines are drafted; and final redline language is presented as suggestions in the document management system, never auto-applied. The platform records every accept, edit, and reject action, and these decisions feed a quarterly process the firm uses to refine the precedent corpus and the firm-tuned prompts.

SUMATOSOFT

Integration and deployment

The platform was deployed inside the Client's Azure tenant to keep all contract content within the firm's existing data perimeter. We integrated with the firm's iManage document management system for contract intake and redline output, and with Active Directory for role-based access aligned to practice groups. A pilot ran with two practice groups over six weeks before firm-wide rollout, with usage metrics and clause-extraction accuracy reviewed weekly with the partner sponsor.

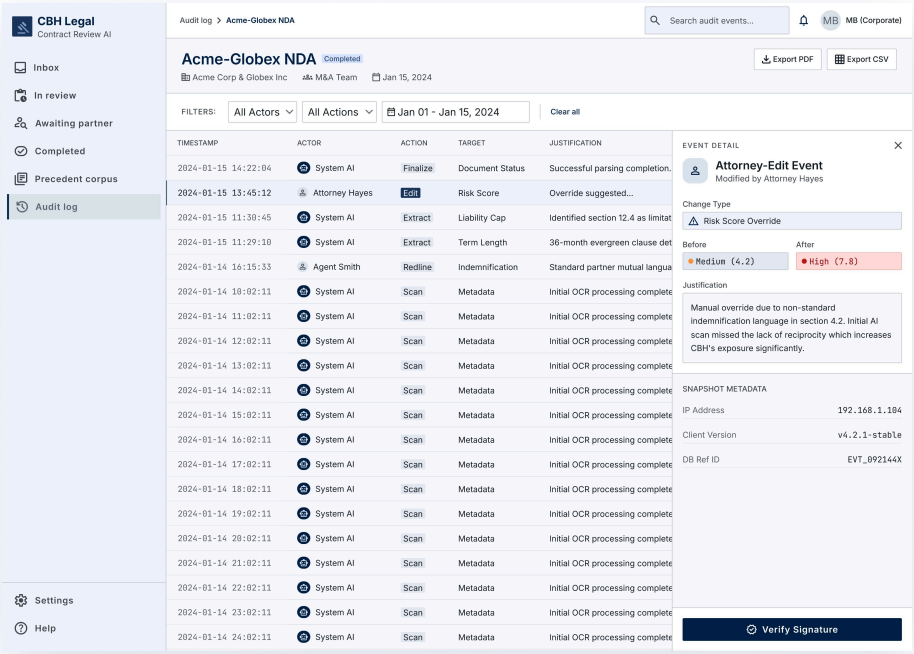

Customer’s benefits

The platform reduced average associate review time on the five covered contract types by 47% and was rolled out to all 80 attorneys across the firm's four practice groups within four months of MVP delivery. Clause coverage consistency, measured against partner-validated checklists, improved from 71% to 94%.

What’s happening with the project right now?

The Client is using the platform across all four practice groups and runs a quarterly precedent-corpus refresh owned by the firm's knowledge management partner, with the next refresh scheduled to add a sixth contract type to the supported set.

sumatosoft.com

info@sumatosoft.com

One Boston Place, Suite 2602 Boston, MA 02108, United States