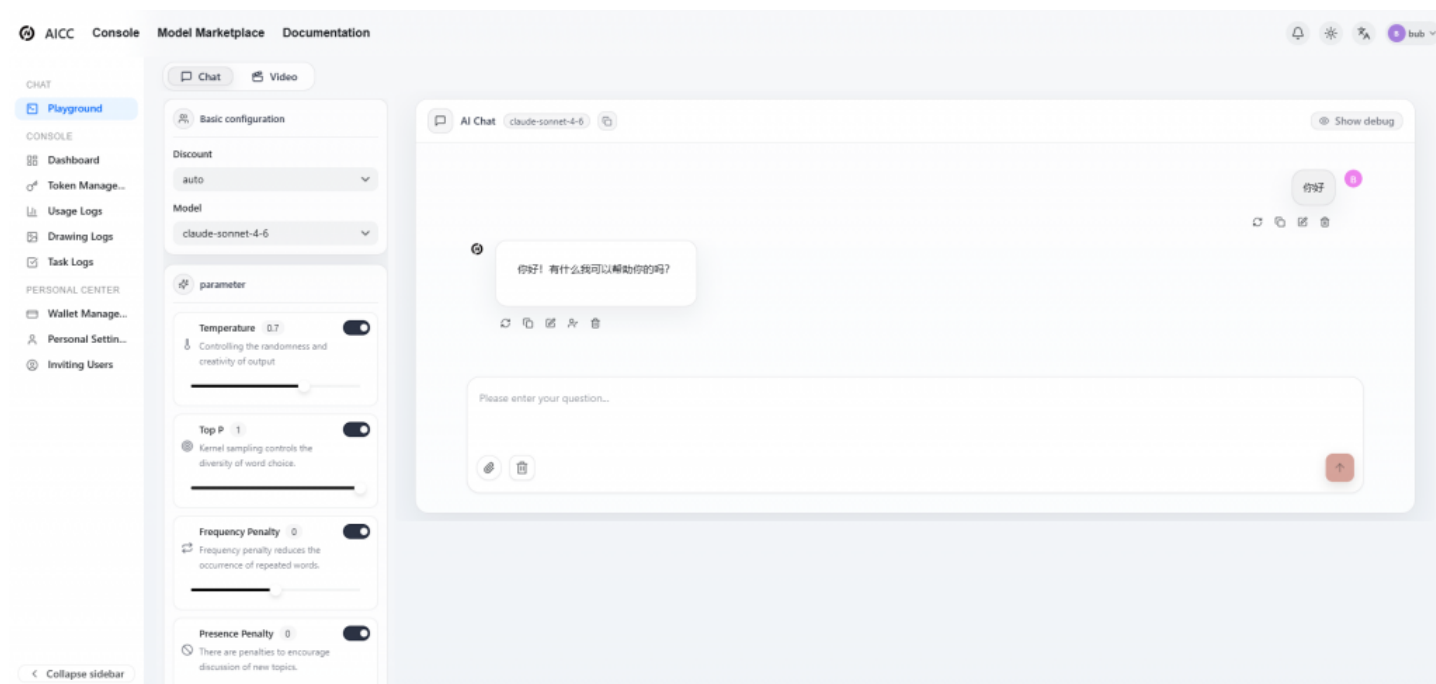


# Legal Tech Startup Automates 70% of Contract Review Workload Using AI.cc Multi-Model API Infrastructure



**Singapore, Singapore May 28, 2026 ([Issuewire.com](https://www.issuewire.com)) - Singapore-based LegalMind AI deploys a five-model routing architecture across Claude Opus 4.7, GPT-5.5, and DeepSeek V4 to cut contract review time from 4.2 hours to 38 minutes per document while reducing AI infrastructure costs by 76%**

AI.cc, the Singapore-based [unified AI API aggregation platform](#), today published a case study detailing how LegalMind AI, a Singapore-based legal technology startup, automated 70% of its contract review workload using AI.cc's multi-model API infrastructure — reducing average contract review time from 4.2 hours to 38 minutes per document and cutting AI infrastructure costs by 76% compared to its previous single-provider deployment.

The case study documents one of the most detailed public accounts of a production multi-model AI deployment in the legal technology sector, providing engineering teams and legal operations professionals with a replicable architecture for AI-assisted contract review at enterprise scale.

LegalMind AI serves mid-market enterprises across Southeast Asia, helping legal teams process commercial contracts, vendor agreements, employment documents, and regulatory filings. Prior to adopting AI.cc's platform, the company was routing all AI workloads through a single frontier model — incurring premium pricing on every step of its document processing pipeline, regardless of task complexity. Monthly AI infrastructure costs had reached a level that was directly constraining the company's ability to price competitively in its target market.

"We were paying frontier model prices for tasks that did not require frontier capability," said LegalMind AI's CTO. "Clause extraction from a standard NDA does not need Claude Opus 4.7. Risk scoring on a

well-structured commercial contract does not need GPT-5.5. We were essentially driving a Formula 1 car to get groceries, and paying for the fuel accordingly."

### The Challenge: Scaling Contract Review Without Scaling Costs

Contract review is a document-intensive, multi-step workflow that maps naturally to agentic AI architecture. A single commercial contract review engagement at LegalMind AI involves eight discrete processing steps: document ingestion and formatting normalization, clause identification and extraction, clause classification by type and risk category, standard clause comparison against approved templates, deviation flagging and risk scoring, regulatory compliance checking against applicable jurisdictions, summary report generation, and human review queue prioritization.

Before adopting AI.cc's platform, all eight steps were processed through a single frontier model. The architecture was simple to implement and produced high-quality output — but the cost structure was unsustainable at scale. Steps one, two, three, and eight — document formatting, clause extraction, classification, and queue prioritization — are structurally simple tasks where a cost-efficient model produces output indistinguishable from a frontier model. Yet LegalMind AI was paying frontier model pricing for all of them.

At the company's Q4 2025 processing volume of 3,400 contracts monthly, averaging 47 pages per contract, the monthly AI infrastructure bill had reached a level that represented 34% of total operating costs — an unsustainable ratio for a growth-stage startup competing on price in a market where incumbent legal service providers have significant cost advantages.

The engineering team evaluated three approaches: negotiating volume discounts directly with its existing frontier model provider, switching entirely to a lower-cost single provider, and migrating to a multi-model architecture on a unified API platform. The first option offered marginal relief. The second required accepting quality reductions on the high-stakes steps where frontier model capability was genuinely necessary. The third offered cost reduction without quality compromise — if the routing architecture could be implemented correctly.

### The Solution: Five-Model Routing Architecture on AI.cc

LegalMind AI's engineering team designed a five-model routing architecture on AI.cc's platform, matching each of the eight contract review steps to the model best suited for its specific requirements.

**Step 1 — Document ingestion and formatting normalization** routes to **Gemini 3.1 Flash** for its native multimodal capability and strong performance on structured document processing. PDF, DOCX, and scanned document inputs are normalized to a consistent structured format for downstream processing. Cost: \$1.00/M input tokens.

**Steps 2 and 3 — Clause identification, extraction, and classification** route to **DeepSeek V4-Flash**. Clause extraction from standard commercial contracts is a pattern-recognition task where a cost-efficient model performs equivalently to frontier models on LegalMind AI's internal quality benchmarks. At \$0.14/M input tokens, routing these high-volume steps to DeepSeek V4-Flash rather than a frontier model represents the single largest cost saving in the architecture. Token volume for extraction and classification represents approximately 40% of the pipeline's total input token consumption.

**Step 4 — Standard clause comparison against approved templates** routes to **Claude Sonnet**

**4.6** for its precise instruction-following and structured output generation. Template comparison requires reliable adherence to comparison rubrics and consistent output formatting across thousands of monthly comparisons — capabilities where Claude Sonnet's instruction-following quality provides measurable advantages over cost-efficiency tier models. Cost: \$3.00/M input tokens.

**Steps 5 and 6 — Deviation flagging, risk scoring, and regulatory compliance checking** route to **Claude Opus 4.7**. These are the steps where frontier reasoning capability directly affects business outcomes. A missed high-risk deviation or incorrect regulatory compliance assessment carries material liability for LegalMind AI's clients. The company's internal evaluation confirmed that Claude Opus 4.7 outperformed every other model tested on these specific tasks, with a 94% agreement rate with senior lawyer review compared to 81% for the next best alternative. Frontier pricing is justified here because the quality differential is real and consequential. Cost: \$5.00/M input tokens.

**Steps 7 and 8 — Summary report generation and queue prioritization** route to **GPT-5.5** and **DeepSeek V4-Flash**, respectively. Summary report generation benefits from GPT-5.5's structured output capabilities and professional prose quality for client-facing deliverables. Queue prioritization — ranking contracts by urgency and complexity for human reviewer assignment — is a classification task routed to DeepSeek V4-Flash.

Implementation: From Single-Provider to Five-Model Architecture in 11 Days

LegalMind AI's engineering team completed the migration from their existing single-provider architecture to the five-model AI.cc deployment in 11 working days — significantly faster than the 6-week timeline the team had originally estimated.

The acceleration was attributable to two factors. First, AI.cc's OpenAI-compatible API formatting meant the team's existing SDK integration required only endpoint and model parameter changes rather than re-engineering. The core API call structure, authentication pattern, and response handling code were unchanged. Second, AI.cc's unified billing and API key management eliminated the need to establish separate vendor relationships, negotiate terms, and integrate billing systems for each of the five models used in the new architecture.

The team used the first three days to run parallel evaluation — processing 200 contracts through both the existing single-model architecture and the proposed five-model architecture, comparing output quality at each step against a ground-truth dataset reviewed by LegalMind AI's senior legal team. The evaluation confirmed quality equivalence or improvement at every step, with the frontier-model steps (risk scoring and compliance checking) showing marginal quality improvement attributed to Claude Opus 4.7's superior performance on legal reasoning tasks compared to the general frontier model previously used.

Days four through eight were spent implementing routing logic, configuring the OpenClaw agent framework for the eight-step workflow, and setting up per-step cost monitoring and quality logging. Days nine through eleven were a staged production rollout, beginning with 10% of live contract volume and scaling to 100% as monitoring confirmed expected behavior.

Results: Six Weeks Post-Deployment

Six weeks after full production deployment, LegalMind AI's outcomes across four measured dimensions exceeded projections.

**Cost reduction:** Monthly AI infrastructure costs fell 76% from the pre-migration baseline — exceeding the 65% reduction the engineering team had projected. The larger-than-expected reduction resulted from the higher proportion of pipeline token volume concentrated in the extraction and classification steps routed to DeepSeek V4-Flash, which consumed more tokens than projected due to longer average contract lengths in the live production distribution versus the evaluation dataset.

**Processing time:** Average contract review time from document submission to completed AI analysis fell from 4.2 hours to 38 minutes — an 85% reduction. The improvement reflects both the efficiency of the parallel processing architecture enabled by the multi-model routing and the elimination of queue bottlenecks that had formed at the single-model endpoint under peak load.

**Automation rate:** 70% of contracts processed post-deployment required no human intervention beyond final sign-off review — up from 41% pre-deployment. The improvement reflects Claude Opus 4.7's superior performance on risk scoring and compliance checking, which reduced the rate of uncertain outputs requiring human escalation.

**Throughput capacity:** Peak processing capacity increased from 180 contracts per day to 640 contracts per day — a 256% increase — without infrastructure changes beyond the API routing migration. The capacity increase reflects the elimination of rate limit constraints that had throttled throughput under the single-provider architecture.

"The business impact was faster than we expected," said LegalMind AI's CTO. "Within three weeks of deployment, we had signed two enterprise contracts that we had previously lost on pricing. The cost structure change made our pricing competitive in a segment we could not reach before."

### Accessing the Full Case Study

The complete LegalMind AI case study, including architecture diagrams, routing configuration details, evaluation methodology, and implementation code examples, is available at [docs.ai.cc/case-studies/legaltech](https://docs.ai.cc/case-studies/legaltech).

Engineering teams interested in replicating the architecture for their own document processing or legal technology workloads can register for a free API key at [www.ai.cc](https://www.ai.cc) and access AI.cc's full model catalog immediately upon registration.

### About AI.cc

AI.cc is a unified AI API aggregation platform headquartered in Singapore, providing developers and enterprises with access to 312 AI models — including GPT-5.5, Claude Opus 4.7, Gemini 3.1 Pro, DeepSeek V4, Llama 4, Qwen 3.6-Plus, and more — through a single OpenAI-compatible API. Additional offerings include the OpenClaw AI agent framework, enterprise SLA plans, AI Translator API, and AI Web Scraping API.

Case study: [docs.ai.cc/case-studies/legaltech](https://docs.ai.cc/case-studies/legaltech) Free API access: [www.ai.cc](https://www.ai.cc) Enterprise plans: [www.ai.cc/enterprise-plans](https://www.ai.cc/enterprise-plans)

## Media Contact

AICC

\*\*\*\*\*@ai.cc

<https://www.ai.cc>

Source : AICC PTE. LTD.

[See on IssueWire](#)