GRAIL

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Methodology You Own

Your Thomson Reuters renewal lands in fourteen months. Harvey just negotiated a data-rights clause into its next enterprise contract. Your senior partner retires in Q4 with the methodology priced into his last three deals still in his head. A briefing for technology and product leaders at mid-size European law firms and compliance advisory practices.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-size legal and compliance technology leaders

DATE

April 2026

EXECUTIVE SUMMARY

Monday 9:15, management committee pre-read

Every CTO, CPO, and Head of Innovation at a mid-size European law firm has had the same Monday. The managing partner has the Thomson Reuters renewal on his desk. Fourteen months until the contract lapses. Eleven percent price increase. Protégé upgrade bundled. Harvey's enterprise team emailed the innovation office Friday about a new data-rights clause in the Q2 contract refresh. The senior partner who ran the firm's largest M&A book for fifteen years retires in Q4. His successor has been named. His methodology sits in three conference-room conversations and one iManage folder. The AI Act Article 4 training obligation lands in August and the compliance officer the firm just hired is asking your office for a system of record.

You have run the technology function for eight years. The pattern used to be simple. Renew the research platforms. Defend the document-management stack. Ship Kira to the corporate team. Keep the partners in Word. The firm was understandable.

Something changed. What used to take twenty paralegals on document review now runs on two senior associates plus a copilot. That compression arrived faster than the review discipline to absorb it. Ninety-five percent of corporate AI deployments produce no measurable P&L impact, and the gap between what the tools can do and what your firm's system actually absorbs is where your seat now lives or loses. Harvey just signed three more Magic Circle firms at enterprise scale. Your senior partner's successor has not been in half of the meetings where the methodology was set. And the CTO of your second-largest corporate client forwarded you a one-line note last month: "We've built a Spellbook deployment in-house. Worth a conversation on scope."

You are not running one technology function. You are running two, and only one of them is on your committee scorecard. One funds what you already run: Thomson Reuters, iManage, Kira, Harvey licences, the KM graph as it sits. The other funds what has to exist by 2028: the data rights you have not renegotiated, the model-agnostic memory store you have not built, the closed-matter corpus you have not structured for agent consumption, the methodology-IP walking out the door with every partner retirement. On one budget line the first wins every renewal cycle. On one scorecard the second does not exist.

This briefing covers what is happening in your sector, what it means for the technology function you run, and the three questions your management committee should answer before the next renewal.

LENS ONE · YOUR INDUSTRY

The legal category just lost the equation that priced the firm

The AI-natives in your category are not experimenting. They are building. Harvey hit a large share of the AmLaw 100 in one year. Legora is embedded in the daily workflow at Mannheimer Swartling and Bird & Bird. Spellbook, EvenUp, CoCounsel, Paxton, Robin AI each own a slice of the drafting, review, and research layer that used to be priced at associate hours. Your corporate clients' general counsel run comparison quotes where your associate rate sits next to a Harvey seat-plus-partner-review structure. They are not asking which produces better drafts. They are asking which pricing model their CFO can defend.

The moat the firm ran on for a decade has moved. The drafting and research layer commoditises fast now. A contract review that used to take a three-associate team two weeks can be run by one senior associate plus Kira plus Harvey in three days, on a modern stack, at a cost the client can measure. Your associate-hour edge was never going to hold. What can hold is the layer underneath the model: the structured methodology your partners encoded into the firm's closed-matter corpus, the data rights you renegotiated before Thomson Reuters absorbed the reference advantage, the integration depth your platform has wired into your client's in-house legal systems. Those compound. They take years to build and they cannot be bought. An AI-native legal platform cannot ship them because they do not have your closed-matter history.

> *Your Harvey deployment this year is on an eighteen-month commoditisation curve. Your methodology-IP encoded as institutional infrastructure is the asset no AI-native can replicate without your closed-matter corpus.*

The industry's speed mismatch is the third force. AI-natives and embedded reference platforms build from scratch or on top of their own dataset with no ethical-wall constraints, no partner compensation politics, no Thomson Reuters contract to answer to inside their own house. Your firm has all three. You cannot pivot at their speed on their terms. What you can do is identify what is genuinely defensible in the firm's methodology, stop defending what is not, and ship the moat they cannot build. The firms that are already moving are quiet about it. The ones that are rolling out Harvey and calling it an AI strategy are the ones the management committee will regret in 2028.

LENS TWO · YOUR FUNCTION

Your technology function is feeling it every day

Your build engine. Harvey and Kira adoption is at seventy percent across corporate. Partners sign off on AI-drafted first passes the associate edited in ten minutes. Nobody is auditing the review depth. Matter cycle time is flat on the dashboard and the firm is shipping drafts with citations it has not verified. The review discipline that worked when associates produced the first pass is training the wrong practice on agent-generated volume now.

Your platform and data. Mannheimer Swartling spent four years on the MSA Innovation Lab foundation that lets Legora and in-house tools actually ship into the partners' hands. You cannot replicate four years in eighteen months. You can buy the drafting-and-research discipline from the tool layer and redirect the compounding capital into the layer no vendor sells: data rights renegotiated across your next Thomson Reuters and LexisNexis contracts, a model-agnostic agent-memory store trained on your closed-matter corpus, an integration depth graph published for your top-20 client in-house legal platforms. The drafting layer commoditises. The methodology layer compounds.

Your product thesis. Your managing partner asked for the firm's three-sentence proprietary advantage last quarter. Your answer was the one from three years ago. Your billable-hour rate card assumes AI efficiency will flow to the firm. It does not. Every hour Harvey saves on the associate side is one hour the client does not buy. That is a pricing-architecture question about where the methodology moat gets built, not a billing-rate one. The joint Pricing Architecture Document with the managing partner and your practice heads needs signing before Harvey signs another mid-market firm into a data-rights clause.

Your R&D bench. The firm stopped hiring juniors in 2024. The routine due-diligence and document-review work that trained first-years into mid-levels is now running on Kira and Harvey. No routine work, no training ground, no senior bench in 2030, no internal candidate for the partner seat in 2032. OpenAI and Anthropic both resumed junior hiring in late 2025 specifically to pair juniors with seniors on judgment-intensive work. The firm that cuts the associate class for this year's margin has a cheaper pyramid in Q4 and an empty bench in five years.

Your committee-level authority. AI-strategy ownership at most mid-size European firms sits with the managing partner in name and nowhere in practice. You are the one writing the Thomson Reuters renewal memo. You are the one the compliance officer is asking for an AI Act system of record. Your seat has the work. It does not yet have the authority. Every month without a written boundary with the managing partner, the AI thesis drifts further from the seat that can defend it.

CAPITAL. VELOCITY. DEFENSIBILITY.

The moat that matters in 2028 is not above the model layer

Your sector is repricing from three sides at once. Harvey and the AI-native platforms absorb the drafting and research layer. Thomson Reuters and LexisNexis embed AI into the reference platforms you already licence. Your corporate clients deploy Spellbook and CoCounsel internally and reduce the scope of what they instruct. Same force, three altitudes. A fourth thing is happening that most mid-size legal CTOs have not named yet.

The moat that matters in 2028 is not above the model layer. It is underneath it, in the methodology your partners carry and your closed-matter files encode. The Harvey deployment you rolled out this quarter is on an eighteen-month commoditisation curve. What does not reset: the data you have rights to, the agent-memory corpus you own, the structured methodology-IP your partners built into closed matters, and the integration depth you published for your top-20 clients. A two-person AI-native cannot build that layer because they do not have your closed-matter history.

Legal firms have lived through this wave before. Before the 1980s, a firm's reference-library advantage was a real moat. Partnership fees paid for the bar-library staff and associates spent weekends there. Westlaw and LexisNexis arrived, and within fifteen years the library advantage was gone. The firms that kept investing in the library paid for a moat that no longer existed. Bloomberg Law absorbed the research-platform layer next. Each wave, the moat moved underneath. Harvey and Protégé are the new Westlaw. The methodology-IP layer is what still sits underneath them.

> *You can tell which camp a firm is in from a single question. Is the technology budget your committee defended at the last renewal a running-cost budget, a compounding methodology-IP budget, or one budget pretending to be one when it is two? Most mid-size firms cannot answer. The ones that can have already made the shift.*

You have twelve to eighteen months. The CTOs and CPOs who claim the authorship of this decision before the next Thomson Reuters renewal author the decade. The ones who sign the standard contract and hand the AI governance file to a compliance officer inherit a seat that has narrowed to renewal negotiation and Harvey licence tracking, and a committee scorecard that stopped mattering the quarter the firm's second-largest client went in-house on Spellbook.

TWO PATHS

What most firms are doing. What the winners are doing instead.

Every legal CTO has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH	THE WINNERS' PATH
Ship Harvey, track licence penetration Seventy-percent adoption. Cycle time flat. Partners sign drafts edited ten minutes. Citations unverified.	**Buy the tool, redesign the review protocol** Eval harness on first-pass output. Review-time floor per stakes tier. Sign-off logged against AI draft.
Sign the Thomson Reuters renewal at the increase Standard contract, no data-rights clause. Protégé bundle accepted. Reference layer absorbs firm workflow.	**Walk in with a data-rights and audit-log redline** Closed-matter corpus stays firm-owned. Integration depth clauses added. Moat investment, not licence line.
Deploy AI efficiency into the billable-hour card Associate bills fewer hours per matter. Revenue falls. Clients take the saving. Margin compresses.	**Stand one fixed-fee practice line on protected P&L** Joint pricing document with the managing partner. Fixed-fee compliance alongside billable book.
Cut the associate class for the Q1 margin win Cheaper pyramid this year. Empty bench by 2031. Lateral-partner hire at a premium.	**Pair juniors with partners plus agents on matter work** Junior plus partner plus Harvey on corporate. Eighteen-month apprenticeship on judgment.

> *Same firm. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the methodology architecture or signs the renewal memo.*

THREE QUESTIONS

Capital Allocation. Build Velocity. Product Defensibility.

Three questions every legal CTO and CPO is tracking. The third is the crux. The first two are how you earn the right to answer it.

01 · R&D CAPITAL ALLOCATION

Is our technology budget one instrument or two?

One funds the existing stack at lower cost. The other builds the methodology-IP moat that survives Harvey, Protégé, and partner retirements. On one renewal cycle the first wins every quarter. On one scorecard the second does not exist. The CTO who walks into the committee with one budget runs the same programme every peer firm is running.

02 · BUILD VELOCITY

Is our review discipline holding on Harvey-generated output, or are we shipping drafts we have not actually read?

Harvey adoption at seventy percent. Cycle time flat on the dashboard. Partners edit first-pass AI drafts for ten minutes and sign. Citations unverified. Mannheimer Swartling spent four years on the platform foundation that lets Legora ship cleanly. You cannot replicate four years in eighteen months. Buy what the tool layer can carry. Own the review bar that makes it safe to ship.

03 · PRODUCT DEFENSIBILITY

What does our firm do that a Harvey plus a good paralegal cannot copy?

Your drafting edge commoditises every eighteen months when the next foundation model rolls. Your closed-matter corpus, your partner methodology, and your data-rights clauses compound. The methodology that priced your senior partner's last three deals sits in his head and in a folder named after him. The moat sits underneath the model layer, not above it. The window to encode it is eighteen months.

FOUR MOVES

What the firms winning this play are actually doing

Four concrete moves any legal or compliance CTO can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture the firm builds on.

1

Build engine: redesign review discipline for Harvey-generated output

Eval harness on first-pass drafting before the partner sees it. Review-time floor per stakes tier. Opinions above threshold require documented review time, not a ten-minute edit. Citation verification as a first-class pipeline. No verified citation, no send. The dashboard is no longer Harvey licence penetration. It is review depth against output volume.

2

Platform and data: build the methodology moat underneath the model layer

Three programmes that stack over eighteen months. Data-rights audit across Thomson Reuters, LexisNexis, and Harvey contracts, with a redline plan tied to the next renewal. A model-agnostic agent-memory store trained on your closed-matter corpus with ethical-wall protocols enforced at ingest. An integration depth graph published for your top-20 client in-house legal teams. The reference layer commoditises. This layer compounds.

3

Product thesis: stand one fixed-fee practice line on protected P&L

Do not migrate the whole book. Keep billable-hour work for matters where stakes demand it. Ship fixed-fee on a parallel P&L for compliance advisory, CSRD work, AI Act governance, standard contract work. Write the cannibalisation trigger and kill criteria in the next partnership cycle. Joint pricing document signed before the first fixed-fee matter closes.

4

R&D bench: rebuild the junior pipeline around legal-engineer hybrids

The 2024 decision to stop hiring associates saves this year and empties the 2031 bench. Reverse it, and reshape it. Juniors in 2026 start on audit, exception-handling, and paired matter work with a named partner plus agent for eighteen months. Add the legal-engineer hybrid role: lawyer with platform fluency, embedded in the technology function, contributing to methodology encoding.

The only firm that joins methodology authorship to working infrastructure

The legal CTO has four places to go. None can do what we do.

The big tech-transformation and legal-tech consultancies ship a standardised AI-governance framework across a sector book. Their cost structure requires reusability. Your firm is idiosyncratic. The framework produces governance theatre you cannot defend to the committee.

The reference-platform vendors (Thomson Reuters, LexisNexis, Wolters Kluwer) have AI-embedded upgrades on the renewal quote. Their business model is embedded AI on the reference layer. Advising a moat that reduces their footprint is impossible for them to say.

The legal AI-native platforms (Harvey, Legora, Spellbook) sell lawyer seats. Their product is measured on licence penetration. Telling you to spend freed capacity on data rights and methodology encoding instead of expanded seats is impossible for them to say.

An in-house legal operations lead plus fractional AI governance splits the seat. Ops is operationally-measured; governance is compliance-measured. One person cannot hold both authorship roles. Two create parallel decisions.

> *The four options you have already looked at. None can run the methodology authorship the seat requires. GRAIL is the fifth.*

WHAT YOU GET	WHY IT MATTERS
Senior advisory that speaks platform-review, not legal-tech deck	Founded by operators who have carried CPO-adjacent seats
Working infrastructure your team uses the same week	Business first, technology second
Custom agents built on your closed-matter corpus, your KM graph	Three weeks from first meeting to working infrastructure
The start of a methodology moat your firm owns	Your team keeps the infrastructure, you own the IP
No vendor lock-in at any phase of the engagement	Every next step earned by what just shipped, never assumed

Three sessions. Working infrastructure from week one.

We start with your technology leadership team because the change cannot outrun the team's understanding of what is possible. Three workshops. Each earns the right to the next.

WORKSHOP 01

How to think

Walk through what is happening in your sector and inside your technology function. Build the first piece of infrastructure in the room: a methodology-IP audit agent that takes your closed-matter corpus and maps what is encodable before the next partner retirement. Your leadership team leaves with a shared view and the first working artefact in hand.

WORKSHOP 02

Your methodology infrastructure

A half-day build. Your Head of KM, Head of Innovation, AI compliance lead, and practice-area champion each leave with five to ten pieces of working infrastructure. Data-rights redline playbook. Closed-matter corpus ingestion prototype. Eval harness for Harvey output. Integration depth graph. Used the same week, not decks.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The committee document that splits the technology budget into two instruments with two hurdle rates. The boundary agreement with the managing partner on who authors the AI thesis. The joint pricing document with the practice heads. The eighteen-month methodology moat build plan.

> *Three weeks. No IT project. No partnership vote needed to start. The infrastructure your technology function is already running on when we finish.*

FIVE QUESTIONS

For your next management committee

The questions where your technology and legal leadership cannot agree on an answer are the ones worth an hour on the agenda. Ask these honestly.

> *Is your technology budget today one instrument or two? If one, what does the committee see that splits running cost from methodology-IP, and what is the kill criterion on each?*

> *Name the AI-native platform in your practice area. What pricing model does it run, and what does the CTO of your second-largest corporate client already have deployed in-house? If you cannot name the top three, that is the first thing the briefing changes.*

> *If your senior partner retires in 2027, how many months would it take to reconstruct the methodology behind the firm's last three major matters from files the junior associate can actually read?*

> *Harvey adoption is at seventy percent. Matter cycle time has not moved. Where did the freed hours go, and is your review bar still the one you want the firm's first-pass drafting trained on?*

> *Your Thomson Reuters renewal is due in fourteen months. Is your redline for data rights, audit logs, and closed-matter corpus exclusions written, or waiting until after the account team presents the standard renewal?*

> *If any of these produces a disagreement in your technology and legal leadership, that is the conversation worth having before the management committee asks.*

JOHAN GRÖNSTEDT · GRAIL · JOHAN@GRAIL.WORKS · GRAIL.WORKS

IN THIS SERIES

There’s a version of this briefing for every seat at the table.

GRAIL has produced industry briefings calibrated to the specific reality of each C-suite role.

CEO

Revenue. Margin. Valuation.

Three numbers your board reads. AI is rewriting all three in your industry.

CFO

Capital. Forecast. Story.

What you allocate this year is the story your board reads next year.

COO

Cheaper. Better. Faster.

The triple threat that used to be a tradeoff is now table stakes.

CHRO

Capacity. Capability. Promise.

The biggest talent transition of a career. Compete on capacity, not headcount.

CPTO

YOU ARE READING

Moat. Speed. Allocation.

The platform moat that survives 2028 is being chosen this year.

CCO

Margin. Pace. Position.

The whole GTM playbook is being rewritten in real time, across every commercial team at once.

Visit grail.works for the full series.

GRAIL

An AI-native management consultancy helping European mid-market leadership teams figure out what fundamentally changes about their business as AI enters their industry.

www.grail.works