



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Stack You Build Your Firm On

Your technology budget is two instruments pretending to be one. Your methodology sits in partner heads and unstructured matter files. Your scorecard measures DMS uptime and document-assembly turnaround, and neither is the 2026 question. A briefing for CTOs, CPOs, and Heads of Innovation at mid-market European professional services firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market
professional
services technology and product
leaders

DATE

April 2026

Thursday 9:15, partnership council Friday

Every technology leader at a mid-sized European professional services firm has had the same Thursday. DMS uptime at ninety-nine-point-eight. Document-assembly turnaround down thirty percent quarter on quarter. Copilot adoption across the partnership tracking ahead of plan. Your Head of Innovation opens the technology committee with the dashboard. He is mid-sentence when your phone buzzes. The Managing Partner forwards an email from the general counsel at the firm's second-biggest client: "We've started a Harvey pilot on transactional-contract review and we'd like fee quotes on the next three deals. We're targeting forty percent of last year's blended rate." Your best second-year associate, the one running the knowledge-automation work stream, resigns that morning to join Legora. The CFO forwards the Unity Advisory Warburg Pincus launch: "Worth discussing before the next partner retreat."

You have run technology and knowledge for this firm for eight years. The pattern used to be simple. Keep the DMS stable. Renew the research-platform contract. Ship document-assembly upgrades. Get Copilot in without the quality-and-risk partner blocking it. The business was understandable.

Something changed. What used to take fifty associates of research and drafting now runs on five partners with copilots. That compression arrived faster than the 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 the firm's partnership actually absorbs is where your seat now lives or loses. An AI-native at ten percent of your leverage economics just landed its first named reference at one of your biggest clients. Your senior partners carry methodology that has never been structured as institutional infrastructure, and two Legora recruiters called your best engineers last month.

You are not running one technology function. You are running two, and only one of them is on your scorecard. One funds what the firm already ships: the DMS, the research-platform renewals, the document-assembly pipeline, the time capture. The other funds what has to exist by 2028: the knowledge graph built from the firm's precedent library, the data rights renegotiated into the next twenty-four months of engagement letters, the proprietary evals run against the firm's own matter corpus, the methodology turned into portable institutional infrastructure. On one hurdle rate the first wins every partnership council. 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 team should answer before the next partnership council.

Professional services just lost the equation that priced the category

The AI-natives in your category have moved past experimentation. They are building, selling, and winning named reference clients in your book. Harvey crossed forty-two percent of the AmLaw 100 in one year. Unity Advisory launched with three hundred million from Warburg Pincus, explicitly structured without a pyramid, priced on outcome rather than on partner-associate leverage. DiligenceSquared is promising ninety percent cost reduction on due diligence. Legora, Spellbook, Robin AI, Darrow each own a slice of what you thought was a workflow moat. Your clients' procurement teams now run comparison quotes where your blended hourly rate sits next to the AI-native's per-outcome number. They are not asking which is the better service. They are asking which pricing model the GC can defend to the audit committee.

The supply-side pressure is only half of it. The demand-side compression is sharper and already arriving. The GCs and CFOs at your top clients are running their own AI on the work they used to send to you. A contract-review engagement that billed at twenty thousand hours last year is running internally this year at six. A transactional-support matter that ran at sixty percent realisation last year is being quoted fixed-fee by a peer firm this year, at a number your associate leverage cannot match. Client procurement has the iron grip. Seventy-three percent prefer outcome-based pricing. Firms retaining time-based pricing grew two-point-one percent against eight-point-seven for value-based adopters. Industry EBITDA is at a five-year low.

Your document assembled this quarter on GPT-5 is on an eighteen-month obsolescence curve. Your firm's structured knowledge asset is what the AI-native cannot replicate without forty years of your matter files.

The third force is the one most partnerships are not naming yet. The first platform-level rewrite in professional services happened between 1980 and 2000. Lexis and Westlaw were not inevitable; by 2000 no firm had a reporter shelf the associates used, and by 2015 the research layer had become a monopoly priced on per-user subscription. Harvey launched in 2023. Forty-two percent of the AmLaw 100 are running Harvey pilots in the second year. That is the second rewrite, and the vendors that won the first one are not the ones winning the second. Your firm has been through this pattern before. You are reading it in your research-licence renewal contract every January. The question is whether your firm builds the layer underneath this rewrite or rents it from whichever vendor wins.

Your technology function is feeling it every day

Your build engine. Document-assembly turnaround is down thirty percent on the dashboard. Partners are drafting first-pass memos with copilots directly in Word. Associates are running research through CoCounsel, Protégé, and two shadow tools the technology committee never sanctioned. Throughput is up. Audit trail is quietly degrading because every partner is using a different stack on different matters. The review bar that worked in 2022 is training the wrong practice on agent-generated volume now.

Your platform and data. Spotify's best developers have not written code since December because Spotify spent four years on platform foundation. Harvey ships to law firms because the platform was purpose-built for matter context from day one. You cannot replicate four years of foundation in eighteen months. You can buy the tool layer and redirect the compounding budget into the layer tool vendors cannot sell: data rights renegotiated across your next twenty-four months of engagement letters, a knowledge graph built from your firm's precedent library, proprietary evals tuned to your firm's own quality standards and run against your own matter corpus. The tool layer commoditises. The structured knowledge layer compounds.

Your product thesis. The partnership council asked for the firm's three-sentence AI thesis last quarter. Your answer was the one from a Big Four consultant a year ago. Your rate card assumes AI-assisted work will price like associate-produced work. It will not. Harvey runs two pricing mechanics. Unity runs outcomes. You run blended hourly. That is a capital-allocation question about where the moat gets built, not a pricing adjustment. Your CFO and Managing Partner need a joint document with you before a client-side GC forces the conversation on their terms.

Your R&D bench. You quietly stopped hiring first-and-second-year technologists in 2024. The routine integration work that used to train a junior into a senior is now running on agents. No routine work, no training ground, no senior bench in 2030. Separately, your best second-year associates are being recruited by Legora and Unity, because the AI-native firms want their judgement on matter structure. The firm that cuts for this year's number has a cheaper function in Q4 and an empty bench in five years.

Your partnership authority. Thirty percent of European AI-strategy ownership sits with the CIO, twenty-three percent with the CDAO, twenty-one percent with the CAIO. You are not on the list. Not because you abandoned the mandate. You have not yet claimed it. The Managing Partner has a Chief AI Officer or Innovation Partner shortlist on the next partnership council agenda. Every month without a visible claim, the thesis drifts further from the seat that can defend it. Your boundary agreement with the Managing Partner lands before the shortlist does, not after.

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

Your sector is repricing from below because AI-natives are shipping outcomes-priced work at ten percent of your leverage economics. Your function is compressing because five partners with copilots now produce what fifty associates did, and the review bar that worked for fifty cannot hold the volume five produce with agents. Same force, two altitudes. But there is a third thing happening that most mid-market professional services CTOs have not named yet.

The moat that matters in 2028 is not above the tool layer. It is underneath it. The memo a partner drafts this quarter on GPT-5 is on an eighteen-month obsolescence curve. The research output of a copilot is commoditised the day it is written. What does not commoditise: the firm's structured knowledge graph, the engagement-letter data rights, the proprietary evals run against your own matter corpus, the integration depth with the top ten client systems. A two-person AI-native cannot build that layer because they do not have forty years of matter files. You do. The window closes in roughly eighteen months, which is the time between now and when the AI-native has enough of its own matter corpus to start building the same layer.

In 1980, Lexis and Westlaw were not inevitable. Senior partners said electronic research would never replace the reporter shelves. By 2000, no firm had a shelf that associates used. By 2015, the research layer had become a monopoly priced on per-user subscription, and every firm in the AmLaw 200 was paying millions a year for infrastructure the firm did not own and could not walk away from. That was the first platform rewrite. Harvey launched in 2023. By 2025, forty-two percent of the AmLaw 100 were running Harvey pilots. That is the second rewrite. The vendors that won the first one are not the ones winning the second. The firms that will read this briefing in five years sit on one side of that pattern or the other.

You can tell which camp a firm is in from a single question. Is the technology budget you defended at the last partnership council a plateau budget, a compounding budget, or one budget pretending to be one when it is two? Most mid-market partnerships cannot answer. The ones that can have already made the shift.

You have twelve to eighteen months. The technology leaders who claim authorship of this decision in Q1 author the decade. The ones who wait until the Chief AI Officer shortlist arrives inherit a seat that has narrowed to DMS uptime and document-assembly throughput, and a scorecard that stopped mattering the quarter the GC at your biggest client opened a Harvey pilot.

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

Every technology leader has had the AI conversation. Most partnerships are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH

Roll out Copilot and CoCounsel, track adoption

Adoption up. Audit trail degrading. Partners routing around the DMS. Knowledge capture falling off a cliff the dashboard does not show.

Route AI budget through IT, sign the requisition

Head of Innovation owns the roadmap within eighteen months. Chief AI Officer arrives in Q3. CTO seat narrows to platform uptime.

Ship AI-assisted work at the blended hourly rate

Client value rises thirty to sixty percent. Realisation flat. A client-side GC forces a forty-percent fee discount on the next three deals.

Cut the junior associate intake for the Q1 partner-pool win

Cheaper leverage ratio this year. Empty senior bench in 2031. Your best seconds get recruited by Legora anyway.

THE WINNERS' PATH

Buy the tool layer, redesign the review bar

Matter-note structure mandatory on AI-assisted output. Audit-trail capture tied to billing. Partner workflow channelled through the DMS, not around it.

Route technology capital through the partnership council

Two budgets, two hurdle rates, two kill criteria. Plateau delivers margin. Compounding delivers the 2028 moat.

Stand one outcomes-priced line on protected P&L

Joint pricing with CFO and Managing Partner. Hourly holds the installed book. Outcomes wins net-new. No cannibalisation.

Pair juniors with seniors on judgement from day one

Eighteen-month apprenticeship on pattern recognition, not five years on routine research the agents absorbed.

Same partnership. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the capital allocation or signs the requisition.

Build Velocity. Product Defensibility. R&D Capital Allocation.

Three questions every professional services technology leader is tracking. The third is the crux. The first two are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our AI adoption shipping production-grade work, or losing audit trail under the dashboard?

Document-assembly turnaround down thirty percent. Copilot adoption ahead of plan. Every partner is using a different stack on different matters and the DMS is emptying. Harvey is purpose-built for matter context. You cannot replicate that in eighteen months. Buy what the tool layer can carry. Own the discipline that makes it ship through the DMS.

02 · PRODUCT DEFENSIBILITY

What does our methodology do that a Harvey at ten percent of our leverage cannot copy?

Your workflow edge commoditises every eighteen months as the next copilot rolls. Your structured knowledge, your engagement-letter data rights, and your client-system integrations compound. Your methodology sits in three senior partners' heads, and two of them retire before 2030. The moat sits underneath the tool layer. The window to build it is eighteen months.

03 · R&D CAPITAL ALLOCATION

Is our technology budget one instrument or two?

Plateau capital runs the existing DMS, document assembly, and research platform at lower unit cost. Compounding capital builds the 2028 moat. On one hurdle rate the first wins every partnership council. On one scorecard the second does not exist. The CTO who walks in with one budget runs the same programme every peer firm is running. The one who walks in with two authors the decade.

What the firms winning this play are actually doing

Four concrete moves any professional services technology leader can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture the firm builds on.

1

Build engine: absorb partner-driven AI usage without the audit trail degrading

Instrument review-bar compliance on AI-drafted deliverables. Matter-note structure mandatory on every AI-assisted output: no structured note, no matter close. Audit-trail capture tied to billing and engagement-letter compliance, not to partner discipline alone. Review-time floor per deliverable above complexity threshold. Your DMS becomes the substrate the AI runs on, not the thing the AI routes around. The dashboard is no longer document-assembly turnaround. It is audit-trail integrity against AI-assisted output volume.

2

Platform and data: build the methodology-as-moat underneath the tool layer

Three concrete programmes that stack over eighteen months. A knowledge graph built from the firm's precedent library, prioritised by matter type rather than by partner politics. Data rights renegotiated across your top-ten-client engagement-letter templates, tied to your next twenty-four months of renewals. Proprietary evals tuned to the firm's own quality standards and run against your matter corpus, not against a vendor's benchmark. The tool layer commoditises. This layer compounds because no AI-native has forty years of your matter files.

3

Product thesis: stand one outcomes-priced line on protected P&L

Do not migrate the whole book. Keep hourly for the installed base that does not want to switch. Stand one outcomes-priced line on a separate P&L for net-new and for clients explicitly asking. Written cannibalisation trigger and kill criteria on day one: if outcomes revenue on an account exceeds hourly revenue, trigger the migration conversation. The engagement-letter template supports both mechanics natively. The joint pricing-architecture document with the CFO and Managing Partner is signed before the first outcomes-priced engagement opens.

4

R&D bench: rebuild the junior pathway around senior-and-agent pairing

Your 2024 decision to cut the junior associate intake saves this year's partner pool. It empties the 2031 senior bench. Reverse it. Juniors in 2026 do not start on routine research or first-pass drafting. The agent layer absorbed that work. They start on audit, exception handling, and paired judgement work with a named senior partner for eighteen months. The pathway compresses from five years to under two because the senior's judgement is now legible in the agent-and-junior review loop. By 2029 you have a mid-tier bench the competition cannot hire because it did not exist anywhere else to hire from.

The only firm that joins methodology authorship to working tools

The technology leader has four places to go. None can do what we do.

The big tech-transformation practices ship a standardised AI-governance framework scaled across a sector book. Their cost structure requires reusability. Your firm's methodology is idiosyncratic. The framework produces governance theatre you cannot defend to the council.

The legal-tech and practice-management platform vendors have a quote on your desk. iManage RAVN, NetDocuments ndMAX, Intapp Walls, Clio Duo, Thomson Reuters CoCounsel. Their product is measured on seat penetration. Telling you to spend compounding capital on the layer they do not sell is structurally impossible.

The AI-native services entrants are not selling to you. Harvey, Legora, Unity Advisory, DiligenceSquared sell directly to the partner buying the work, around the CTO. Your alternative is a build-versus-partner-versus-compete decision that sits on your desk.

An in-house or fractional Chief AI Officer splits the seat. Strategy CAIO is revenue-measured; Platform CAIO is deployment-measured. Two-thirds of newly-appointed CAIOs spend year one compiling inventories instead of setting direction.

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

WHAT YOU GET

Senior advisory that speaks partnership-council, not transformation-deck

Working tools your team uses the same week

Custom agents built on your matter corpus, your precedents, your engagement letters

The start of a methodology-as-infrastructure moat your firm owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried senior partnership-facing seats

Business first, technology second

Three weeks from first meeting to working tools

Your firm keeps the tools, you own the IP

Every next step earned by what just shipped, never assumed

Three sessions with your technology leadership. Working tools from week one.

We start with your technology leadership team because the change cannot outrun the team's personal understanding of what is possible. Three workshops. Each delivers standalone value. 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 tooling in the room: a capital-allocation agent that takes your real technology budget, splits it two ways, and runs the alternatives analysis against your named AI-native competitors. Your leadership team leaves with a shared view and the first working artefact already in hand.

WORKSHOP 02

Your firm's AI tooling

A half-day build. Your Head of Knowledge, Head of Product, Head of Engineering, and Head of Innovation each leave with five to ten pieces of working tooling tuned to their function. Review-bar audit agent. Matter-note structure scaffolding. Data-rights renegotiation playbook. Integration depth graph prototype. Tools your team is using the same week, not decks about tools.

WORKSHOP 03

Ready to run

Refined tools based on real use. The partnership-council document that splits the technology budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the Managing Partner on who authors the AI thesis. The joint pricing-architecture document with your CFO and senior partners. The eighteen-month moat build plan.

Three weeks. No IT project. No partnership council approval needed to start. The tools your technology function is already running on when we finish.

For your next partnership council

The questions where your technology 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 partnership council see that splits plateau from compounding, and what is the kill criterion on each?

Name the AI-native in your category. What is their pricing mechanic, and at what percent of your leverage economics are they quoting? If you cannot name the top three, that is the first thing this briefing changes.

If your two most senior partners retired tomorrow, how many months would it take to reconstruct the methodology behind your firm's top-five service lines from what is actually structured in the DMS?

Copilot and CoCounsel adoption is ahead of plan. Audit-trail integrity is sliding. Partners are routing around the DMS. Where did the captured knowledge go, and is your review bar still the one the quality-and-risk partner signed off on?

Your Managing Partner has a Chief AI Officer or Innovation Partner shortlist on the next council agenda. Is your boundary agreement written, or is it waiting until after the shortlist lands?

If any of these produces a disagreement in your technology leadership, that is the conversation worth having before the partnership council asks.

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