



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Platform You Own

Your regulatory-knowledge platform sits in four senior specialists' heads. RegTech vendors are shipping AI modules that absorb the methodology layer you bill against. Your scorecard measures copilot adoption and monitoring-feed coverage, and neither is the 2026 question. A briefing for CTOs, CPOs, and Heads of Platform at mid-market European regulatory and compliance advisory firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market regulatory
and compliance advisory platform
leaders

DATE

April 2026

Monday 9:15, platform review

Every CTO at a mid-sized European regulatory and compliance advisory firm has had the same Monday. Copilot adoption across the practice at sixty-five percent. Specialist time per CSRD gap analysis compressed from six weeks to two. Your VP Platform opens the review with the dashboard. He is mid-sentence when your phone buzzes. Email from the CEO: "Moody's ships their AI Act module in June. What is our response?" Your senior AI Act specialist mentioned retirement plans Friday. Your CFO forwarded an Advensse press release: "They're positioning RegTech revenue in IK's data room. Worth discussing before Q2."

You have run this function for seven years. The pattern used to be simple. Build the monitoring infrastructure. Ship the copilots. Hold the review bar. Meet the partnership on monitoring coverage and adoption. The business was understandable.

Something changed. What used to take a senior specialist six weeks now runs in two with an agent-augmented junior. 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 tool capacity and review discipline is where your seat now lives or loses. A RegTech vendor with the same regulatory content just quoted your largest financial-services client at a fraction of your mandate fee. Your senior specialists carry enforcement-culture intelligence that has never been written down, and two of them retire inside three years.

You are not running one platform. You are running two, and only one is on your scorecard. One runs the monitoring and drafting copilots you bill against today. The other has to exist by 2028: the enforcement-culture graph you have not built, the response-playbook library you have not encoded, the client-specific memory that outlasts senior retirements. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist.

This briefing covers what is happening in your sector, what it means for the platform you run, and the three questions your team should answer before the next partnership meeting.

Regulatory advisory just lost the methodology moat

The RegTech vendors in your category are not experimenting. They are building. Moody's, Thomson Reuters Regulatory Intelligence, MetricStream, Wolters Kluwer, Diligent each shipped AI modules in 2025 that absorb the methodology layer your senior specialists bill against. They have the regulations. They now have gap-analysis templates, obligation-mapping graphs across AI Act / CSRD / DORA / MiCA, and response-drafting agents. A financial-services client running procurement quotes your firm against a RegTech subscription at a fraction of your mandate fee. They are not asking which is the better advisor. They are asking which cost structure their CFO can defend through the next CSRD cycle.

The moat your firm ran on for twenty years has moved. The methodology layer commoditises fast now. A capability that used to take a senior specialist six weeks to deliver can be templated by a RegTech vendor in a quarter and priced on a per-mandate basis the client's procurement team will choose by default. Your methodology edge was never going to hold. What can hold is the layer underneath the regulations: the accumulated understanding of how the Swedish Finansinspektionen actually sequences inspections, how the Danish Finanstilsynet scores AI Act literacy responses, how BaFin behaves on DORA ICT risk. Those are enforcement-culture patterns. They take years to build. A RegTech vendor cannot ship them because they do not have the mandate base that generated them.

Your gap-analysis template on GPT-5 is on an eighteen-month obsolescence curve. Your enforcement-culture intelligence across forty prior mandates is the asset the RegTech vendor cannot replicate without your client base.

The industry's demand-floor paradox is the third force. CSRD, DORA, and AI Act mandates are non-discretionary. The wave is real and growing. That demand floor protects revenue for the next three years even if the firm does nothing. It does not protect margin. The firms that expand headcount proportionally into the wave build cost structures that competitors with AI-augmented delivery will undercut by 2028. The firms that are already moving are quiet about it. The ones that are shipping copilots and calling it an AI strategy are the ones the partnership will restructure when the wave ends.

Your platform function is feeling it every day

Your build engine. Copilot adoption at sixty-five percent. Specialist hours per mandate down forty percent. Practice heads pushing fifty-page gap analyses through review with minutes of senior attention. Foundation cases miss. Twelve months in, your team cannot execute the novel AI Act interpretations a regulator audit will demand. The review discipline that worked on human-drafted documents is training the wrong practice on agent-generated volume.

Your platform and data. The regulatory-knowledge platform your firm bills against sits in four senior specialists' heads. Two retire within three years. Monitoring feeds half-automated. Obligation-mapping corpus is a shared drive. You cannot buy what Moody's is building because Moody's is building the regulations side, and your moat was never the regulations. Buy the drafting infrastructure from the tool layer. Redirect compounding capital into what the RegTech vendor cannot sell: the enforcement-culture graph encoded from forty prior engagements, a response-playbook library keyed to regulator and jurisdiction, a model-agnostic memory store portable across model swaps.

Your product thesis. Your firm has one revenue model: hours billed against senior judgment. Your rate card assumes AI-augmented delivery prices like manual delivery. It does not. RegTech vendors run subscription pricing that compounds with client count. You run mandate pricing that compresses with AI augmentation. That is a capital-allocation question about whether the platform becomes a product line. Your practice heads and CFO need a joint document with you before the next RegTech quote closes at a client you thought was yours.

Your R&D bench. Your senior AI Act and CSRD specialists are carrying the firm. Two are five years from retirement. Your junior pipeline went quiet in 2024 because AI-augmented delivery lets one senior produce what three did. No juniors, no training ground, no mid-tier bench in 2031, no practice-head successor. The firms that cut juniors for this year's margin have a cheaper function in Q4 and an empty succession pipeline.

Your board-level authority. Thirty percent of European AI-strategy ownership sits with CIO, twenty-three percent CDAO, twenty-one percent CAIO. You are not on the list. Your CEO has a CAIO shortlist. Your practice heads are running their own AI initiatives without your platform discipline. Every month without a visible claim, capital-allocation authority drifts further from the seat that can defend it. Your Q1 boundary agreement with the CEO and practice heads lands before the CAIO shortlist, not after.

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

Your sector is repricing from below because RegTech vendors ship content platforms at subscription economics your mandate fees cannot match. Your function is compressing because one senior plus an agent now produces what three specialists did. Same force, two altitudes. There is a third thing most mid-market regulatory advisory CTOs have not named.

The moat in 2028 is not above the model layer. It is underneath it. The gap-analysis template you ship on GPT-5 is on an eighteen-month obsolescence curve. What does not commoditise: the enforcement-culture intelligence across forty prior mandates, the response-playbook library your senior specialists carry, the client-specific memory your platform stores, the obligation-mapping corpus you maintain. A Moody's module cannot build this because it does not have the mandate base that generated the signal. You do. The window closes in roughly eighteen months.

In the early 1990s, dominant legal-research firms owned curated case-law libraries edited by specialists who knew which decisions were load-bearing in which jurisdictions. Their CTO's scorecard was database coverage, editor productivity, renewals. Every metric was green. Bloomberg Law entered in 2009 with structured data and AI-assisted search. By 2024 the editorial moat had compressed. The 2009 scorecard never turned red. It just stopped being the one the market was scoring. The incumbent lost because the editorial judgment never became a product. Someone else built the platform around the knowledge.

You can tell which camp a firm is in from a single question. Is the platform budget you defended at the last partnership meeting a plateau budget, a compounding budget, or one pretending to be one when it is two? The firms that can answer have already started encoding.

You have twelve to eighteen months. The CTOs who claim authorship in Q1 author the decade. The ones who wait inherit a seat narrowed to monitoring uptime and copilot adoption.

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

Every CTO at a regulatory advisory firm has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH

Ship copilots, track adoption

Sixty-five percent adoption. Drafts through minimal review. Enforcement patterns missed.

Run platform budget through IT

CIO owns roadmap in eighteen months. CAIO in Q3. CTO narrows to monitoring uptime.

Bill hours against the wave

Revenue fine three years. Wave ends. Platform never built. Firm repriced when RegTech has the footprint.

Cut juniors for Q1 margin

Cheaper practice now. Empty succession bench by 2031. External hire premium the firm cannot absorb.

THE WINNERS' PATH

Buy the drafting layer, redesign review

Senior-review floor. Sign-off on novel interpretations. Eval harness on golden-set cases.

Route platform capital through M&A hurdle

Two budgets, two hurdle rates, two kill criteria. Plateau for margin. Compounding for 2028.

Stand one subscription line alongside mandates

Parallel P&L. Joint pricing with practice heads and CFO. Captures non-advisory revenue.

Pair juniors with seniors and agents

Juniors annotate playbook entries from day one. Career curve compresses ten years to three.

Same practice team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the platform-allocation decision or signs the requisition.

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

Three questions every regulatory advisory CTO is tracking. The second is the crux. The first and third are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our specialist throughput shipping regulator-ready output, or drafts that miss the enforcement pattern?

Copilot adoption at sixty-five percent. Specialist hours per mandate compressed forty percent. Practice heads pushing gap analyses through review with minutes of senior attention. The review discipline that worked on human-drafted work is training the wrong practice on agent volume. Buy what the tool layer can carry. Own the discipline that makes it ship past an audit.

02 · PRODUCT DEFENSIBILITY

What does our platform hold that a Moody's or Thomson Reuters module cannot copy?

The RegTech vendors have the regulations. They do not have how Finansinspektionen sequences inspections, how BaFin scores AI Act literacy, or the response playbooks from your last forty mandates. Your moat sits in four senior specialists' heads. Two retire inside three years. Encode it, or watch the knowledge walk out the door the quarter the vendor lands your largest client.

03 · R&D CAPITAL ALLOCATION

Is our platform budget one instrument or two?

Plateau capital runs the monitoring feeds and copilots at lower unit cost. Compounding capital builds the 2028 platform: the enforcement-culture graph, the response-playbook library, the obligation-mapping corpus, the client-portal as product line. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist.

What the firms winning this play are actually doing

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

1**Build engine: absorb agent throughput without the review bar collapsing**

Instrument senior-review time per mandate as a floor. Any gap analysis above a novelty threshold requires named senior sign-off and thirty minutes of documented review. Mandatory annotation on every agent-drafted interpretation. Evaluation harness as first-class infrastructure: golden-set enforcement cases, rubric scored against platform output. The dashboard is no longer copilot adoption. It is review depth against agent volume.

2**Platform and data: build the regulatory-knowledge platform underneath the model layer**

Five components stacking over eighteen months. An enforcement-culture graph across Finansinspektionen, Finanstilsynet, BaFin, DNB, and Datatilsynet, annotated by senior specialists. A response-playbook library keyed to regulator, jurisdiction, and client profile. An obligation-mapping corpus maintained as regulations change across AI Act, CSRD, DORA, MiCA. A client-specific memory store portable across model swaps. And the client-portal, positioned as a subscription product line.

3**Product thesis: stand one platform-subscription line alongside mandate fees**

Do not migrate the whole book. Keep hours-billed for complex novel mandates. Ship platform subscription for ongoing monitoring, obligation tracking, and playbook access. Separate P&L. Write the cannibalisation trigger in Q1: if subscription revenue on an account exceeds mandate revenue, trigger the architecture conversation. The joint Pricing Architecture Document with practice heads and CFO is signed before the first subscription deal closes.

4**R&D bench: encode senior-specialist knowledge before the retirement wave**

Your 2024 decision to stop hiring juniors saves Q1. It empties the 2031 senior bench. Reverse it. Every senior AI Act and CSRD specialist spends one day per week with a named junior plus an agent, producing annotated playbook entries from active mandates. Career curve compresses ten years to three because senior judgment becomes legible in the agent-and-junior review loop. By 2029 the firm has a mid-tier bench the competition cannot hire from.

The only firm that joins platform authorship to working infrastructure

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

The big risk-and-regulatory practices (Deloitte, PwC, KPMG, EY FSRM) ship a standardised framework scaled across a sector book. Their cost structure requires reusability. Your seat is idiosyncratic. A reusable framework on an idiosyncratic seat produces governance theatre.

The RegTech platform vendors (Moody's, Thomson Reuters, MetricStream, Wolters Kluwer, Diligent) have a quote on your desk. They sell regulations and templates. Telling you to encode your own enforcement-culture intelligence instead is impossible for them to say.

AI-native regulatory-adjacent vendors (Harvey, Luminance, Robin AI) compete with your firm. Their product architecture prevents honest advice about where your moat is defensible and where it is already lost.

An in-house or fractional CAIO splits the seat. Strategy CAIO revenue-measured, Platform CAIO velocity-measured. One person cannot hold both. Two-thirds of new CAIOs spend year one compiling inventories instead of setting direction.

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

WHAT YOU GET

Senior advisory that speaks platform-review, not transformation-deck

Working infrastructure your team uses the same week

Custom agents built on your mandates, your clients, your playbooks

The start of a regulatory-knowledge platform your firm owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried C-suite platform seats

Business first, technology second

Three weeks from first meeting to working infrastructure

Your team keeps the infrastructure, you own the IP

Every next step earned by what just shipped, never assumed

Three sessions with your platform leadership. Working infrastructure from week one.

We start with your platform 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 platform function. Build the first piece of infrastructure in the room: a platform-allocation agent that takes your real platform budget, splits it two ways, and runs the alternatives analysis against the named RegTech vendors quoting your clients. Your leadership team leaves with a shared view and the first working artefact already in hand.

WORKSHOP 02

Your platform infrastructure

A half-day build. Your Head of Platform, Head of Data, practice heads, and senior specialists each leave with five to ten pieces of working infrastructure tuned to their function. Enforcement-culture annotation agent. Playbook-encoding scaffolding. Obligation-mapping prototype. Client-memory store skeleton. Infrastructure your team is using the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The partnership document that splits the platform budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the CEO on who authors the platform thesis. The joint Pricing Architecture Document with your practice heads and CFO. The eighteen-month platform build plan.

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

For your next platform review

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

Is your platform budget today one instrument or two? If one, what does the partnership see that splits plateau from compounding, and what is the kill criterion on each?

Name the RegTech vendor quoting your top-ten financial-services clients on the AI Act and CSRD modules. What is their pricing mechanic, and at what percent of your mandate fee? If you cannot name the top three, that is the first thing the briefing changes.

If your senior AI Act specialist or your senior CSRD lead retired tomorrow, how many months would it take to reconstruct the enforcement-culture intelligence behind your last forty mandates?

Copilot adoption is at sixty-five percent. Specialist hours per mandate are down forty percent. Where did the freed hours go, and is your senior-review discipline still the one you want the next agent wave trained on?

Heidrick's February 2026 data: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO on the AI-strategy-ownership list, zero percent CTO-CPO. Is your Q1 boundary agreement with the CEO and the practice heads written, or is it waiting until after the CAIO shortlist lands?

If any of these produces a disagreement in your platform leadership, that is the conversation worth having before the partnership 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