

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in Your Capital Story

Your next €10M of capital. The renewal number you sign. The thesis your board can recite back. Three questions that look separate, but are actually one. A briefing for CFOs at mid-market European research and intelligence firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market
research & intelligence CFOs

DATE

April 2026

Thursday morning, board pack due next Thursday

Every CFO at a mid-sized European research or intelligence firm has had the same Thursday morning. Four things on the screen at 8:30. Q3 syndicate renewal printed at 84% against a 92% plan, the second consecutive quarter below. A top-five client cancelled its subscription eleven days ago citing its own internal AI capability; three top-twenty renewals have gone quiet since. Your CTO wants €2M next quarter to rebuild the delivery platform on a foundation-model architecture, and his one-page says "AI infrastructure investment." Your sponsor's partner forwarded the Bloomberg "AI is eating research" story yesterday with a single line: "curious to hear your thinking on the migration plan before Monday."

You have signed twelve renewal forecasts for this board. The pattern used to be simple. Hold the rate card, hold the attach, grow into the plan. Close the quarter clean. Keep the auditor happy. The business was understandable.

Something changed. Your renewal is drifting and your team cannot reconstruct the account-level story. Your AI capital is running through the IT budget as a productivity line. Your clients are quoting their own AI capability in the procurement conversation, your rate card is per-seat per-syndicate, and the comparison sits badly. Your Controllers are closing faster than they ever have, and the audit committee is asking about deferred-revenue traceability on the custom-inquiry book. The close is clean. The story underneath it is not.

Your production economics are collapsing. The intelligence asset under this firm has never been more valuable. The job has not changed in fifty years. The bar has risen on every measure of doing it, simultaneously, and the contradictions you have been navigating since you took the seat are suddenly re-openable. The CFOs who treat them as engineering problems run the offensive play. The ones who treat them as permanent trade-offs lose capital-allocation authority to the CTO within eighteen months.

This briefing covers what is happening in your sector, what it means for the finance function you run, and the three questions your team should answer before the next board pack.

Mid-market research just lost the production model that priced the sector for a decade

Every AI capability your clients adopt raises their ability to run secondary research, synthesise public sources, and draft structured analysis at near-zero marginal cost. The part of your work your clients used to outsource because it was too expensive to do in-house is the part they now run themselves on Friday afternoon. Procurement teams at your best accounts are citing their own AI capability in the renewal conversation. Your per-seat syndicate number sits next to their internal-build cost, and the comparison is unflattering. AI-native research entrants are entering your sectors with dramatically lower cost structures and three pricing mechanics against your one. Across your installed book the mechanical difference is four to eight points of renewal a year, and it is already showing up in the quarterly prints.

Your competitors are not waiting. Bisnode's path into Dun & Bradstreet is the sector's future shape, not an exception. Standalone research firms are either absorbed into data-platform companies or left competing against them. The firms moving first are buying AI-native research entrants as tuck-ins rather than trying to rebuild themselves. Some are migrating pricing on their custom-inquiry and advisory books first and the syndicate subscription second. A small set of the more patient ones are systematising fifteen years of accumulated client work, sector pattern libraries, and methodology records into proprietary data assets no AI-native entrant can replicate in eighteen months. They are the quiet compounders. The rest of the sector is shipping AI features into the delivery surface at no incremental price and wondering why renewal is compressing.

Your production economics are collapsing. The intelligence asset under this firm has never been more valuable. That is not a product-marketing problem. It is a capital-allocation decision the finance seat has not yet claimed.

The sector is repricing underneath everyone at the same time. GPT, Perplexity, and Gemini synthesise public research competently enough for most client decisions that used to pay syndicate rates. The 6-12 month window is not theoretical. It is the horizon inside which pricing conversations reset around AI-generated alternatives. The firms that treat this as a client-success problem rather than a capital-allocation-and-moat problem are walking into a board conversation they will not win.

Your capital, your forecast, and your narrative are feeling it every day

Your capital allocation. Your CTO's €2M AI-infrastructure request is sitting in the IT budget, routed through the same mechanism you use to sign off on laptops and software licences. It is not a tooling request. It is the largest single capital-allocation decision you will make this decade, and it deserves the same discipline you bring to M&A. The finance seats that ship this as a capital line to the board, with alternatives analysis and an ROIC hurdle, keep the authorship for the next ten years. The ones that sign the requisition hand the authorship to the CIO, the CTO, and the Head of Research by default.

Your renewal forecast. The team built the annual plan on a 92% renewal and a 108% attach assumption. The Q3 print came in at 84% and 96%. Your Head of FP&A cannot reconstruct the account-level story because the churn signal lives in analyst CS notes, the attach signal lives in platform usage, and the pricing-pressure signal lives in procurement benchmarks her team has never had access to. The plan is not wrong because the forecast was aggressive. It is exposed because the data architecture underneath it was built for quarterly reconciliation, not for continuous reconstruction of the renewal story.

Your pricing architecture. The board voted on a pricing strategy that assumed syndicate subscription was the defensible revenue. It is the most exposed line in the business. Your AI-native competitors are on three pricing mechanics. Your custom-inquiry book could be priced per-question or per-insight today; it is being sold at blended day rates. Your advisory retainers could be outcome-linked; they are not. Pricing architecture is not a commercial-director question. It is a capital-allocation question about where the defensible revenue comes from in three years.

Your analyst bench and Controller pathway. The secondary research, literature review, benchmarking, and first-draft synthesis work that used to train a junior analyst into a Senior Analyst into a Research Director are now running on agents. The same is true for the journal-entry and reconciliation work that trained your junior controllers. The Q1 headcount win is real. The 2031 senior-analyst and VP Finance bench is empty. The firm that cuts the junior layer without rebuilding the entry pathway has a cheaper quarter this year and no internal candidate for the seat that matters in five.

Your economic thesis. Your board has asked for the strategy-day preview. Your sponsor's partner has asked for the migration-plan note. Your chair has asked how the firm stays ahead of client-internal AI. Three audiences, three different documents, each drafted by a different team. None of them are anchored in the same data. The CFOs pulling ahead write the thesis once, anchor it in the proprietary-data asset, and let the three audiences draw from the same architecture. The ones behind them are editing three decks and hoping the numbers match.

You do not have three scorecard problems. You have one.

Your sector is repricing from below because AI-native entrants bring three pricing mechanics and direct-to-client substitution pressure, and the rest of the market is on one syndicate price. Your finance function is stretched because the renewal forecast, the capital-allocation authority, and the economic thesis are each being managed as a separate programme, each drawing on a different data layer, each drafted by a different team. Same force, two altitudes. But there is a third thing happening that most research-firm CFOs have not named yet, and it is the one that matters most.

Capital allocation without a credible forecast is guessing. A credible forecast without a narrative is scorekeeping. A narrative without a data architecture underneath it is marketing. The finance seat that holds all three holds one integrated asset, not three separate programmes. Every capital-allocation question, every forecast-credibility gap, every narrative weakness is a symptom of the same missing layer: the proprietary-data infrastructure built from fifteen years of client work, sector pattern libraries, and methodology records, feeding all three from the same data, the same governance, the same discipline.

This is the Reuters-Bloomberg decision of this decade. Every CFO in your cohort is making it right now, whether they see it or not. The firms that built proprietary data infrastructure in 1989-1991 compounded for thirty years. The firms that licensed the new substrate became resellers the market does not remember. The shape is the same. The asset under decision is different. The cost of getting it wrong in 2026 is compressed into eighteen months instead of five years.

You can tell which camp a firm is in from a single question. Can your team answer any material board question by walking the same data spine, in the same meeting, with the same numbers? Most mid-market research firms cannot. The ones that can have already made the shift, and they are the ones the board is already promoting.

You have twelve to eighteen months. The CFOs who own this decision personally run the offensive play. The ones who do not become the CFO the board replaces when the successor walks in with a thesis.

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

Every research-firm CFO has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones. Default looks cleaner for three quarters. Winners compound from month fifteen.

THE DEFAULT PATH

Route AI capital through the IT budget

Sign the CTO's requisition. Productivity line in the annual plan. CIO and Head of Research own the roadmap within eighteen months.

Hold syndicate pricing, ship AI features free

Client-facing AI capability rises. Syndicate renewal drifts four to eight points below plan. Custom-inquiry book sold at blended day rates.

Run renewal on a rolling spreadsheet

Quarterly cadence. Three scenarios take a week. Below-plan quarters produce no account-level story until the post-mortem.

Cut the analyst and Controller layer for the Q1 win

Headcount down fifteen to twenty-five percent. Empty succession bench in 2031. External senior hire the year after.

THE WINNERS' PATH

Route AI capital through the M&A hurdle

Alternatives on one page. Proprietary-data build, AI-native rebuild, tuck-in, working capital, buyback. ROIC-scored.

Migrate pricing to three mechanics

Multi-dimensional on AI-augmented tiers. Per-insight output pricing on custom inquiry. Outcome-share on advisory. Value converts to margin.

Rebuild the forecast as a live driver graph

Ten scenarios in an afternoon. Below-plan quarters produce the account-level story in the same meeting.

Redesign the junior pathway around judgment

New hires pair with agents on sector pattern work and capital exposure. Bench rebuilds faster than it would have grown.

The default path shows a cleaner Q1. The winners' path has visible friction for three quarters and compounds from month fifteen. The difference is whether the CFO is authoring the capital or administering it.

THREE QUESTIONS

Capital Allocation. Forecast Credibility. Value-Creation Narrative.

Three questions every research-firm CFO is tracking. The first is the crux. The other two are how you earn the right to answer it.

01 · CAPITAL ALLOCATION

Where does the next €10M actually compound?

Proprietary-data systematisation? AI-native delivery platform? Tuck-in acquisition of an AI-native research entrant? Working-capital extension as custom cycles lengthen? Share buyback? Today's default: the CTO asks, the IT budget funds, the requisition gets signed, and the capital-allocation authority slides away from the seat. The question is not which tool. The question is whose decision this is.

02 · FORECAST CREDIBILITY

Why did the number print below plan, and what is the account-level story?

Your board voted on 92% renewal. Your firm is printing 84%. Your Head of FP&A cannot walk the audit committee through the account-level story because the churn signal sits in analyst CS notes, the attach signal sits in platform usage, and the pricing-pressure signal sits in procurement benchmarks the finance team has never seen. Three scenarios used to take a week. They now run in an afternoon. The bar has risen on what credible means.

03 · VALUE-CREATION NARRATIVE

What does this firm do that compounds faster than peers?

Your board wants the strategy-day preview. Your sponsor wants the migration-plan note. Your chair wants the how-we-stay-ahead answer. Three audiences, three documents, three teams, and the numbers underneath them do not match. The CFOs pulling ahead write the thesis once, anchor it in the proprietary-data asset, and refresh it continuously from the same data that runs the forecast.

What the firms winning this play are actually doing

Four concrete moves any research-firm CFO can start this quarter. Each addresses one lever. Pulled together, they become one capital-intelligence asset the firm runs on.

1

Capital: route the AI-infrastructure decision through the M&A hurdle, not the IT requisition

Your CTO's €2M request is a capital-allocation decision. Treat it like one. Present the alternatives on one page: proprietary-data systematisation, AI-native delivery rebuild, tuck-in of an AI-native research entrant, working-capital extension, buyback. ROIC-scored against the same five-year horizon you use for M&A. Name what you are choosing not to fund to protect the thing you are funding. That is the sentence that keeps authorship at the seat.

2

Forecast: rebuild the renewal forecast as a live driver graph, not a rolling spreadsheet

Your renewal model should draw from live analyst CS notes, live platform-usage signals, live procurement-benchmark data, and live competitive-intelligence signals on AI-native entrants, as one graph rather than five spreadsheets stitched together. Three scenarios compress from a week to an afternoon. Below-plan quarters produce the account-level story in the same board meeting. The forecast stops being a deliverable on a cadence and becomes a live artefact the CFO invests in continuously.

3

Pricing: migrate from one syndicate mechanic to three

Multi-dimensional pricing on AI-augmented syndicate tiers. Per-insight or per-question output pricing on the custom-inquiry book where the value is now the interpretation, not the information. Outcome-share on the advisory book where research drives a measurable client decision. Syndicate subscription holds the base. The three new mechanics convert AI-augmented value to margin on the accounts where AI-native entrants are already quoting. Finance authors the migration plan. Research builds against it.

4

Narrative: write the economic thesis once, refresh continuously

Three sentences on one page. What does this firm do that compounds faster than peers? What capital gets starved to protect it? What gets fed? Anchor the thesis in the proprietary-data asset that fifteen years of client work has produced. Refresh continuously off the same driver graph that runs the forecast. Sponsor-monthly excerpts the thesis. Audit-committee prep excerpts the thesis. Strategy-day deck excerpts the thesis. Three audiences, one artefact, four registers.

The only firm that joins capital authorship to working infrastructure

The research-firm CFO has four places to go. None can do what we do.

The big finance transformation practices will sell a twelve-to-twenty-four-month close-acceleration programme staffed by a partner and analysts billing hours. Their audit arm constrains their governance advice. Their billable-day model requires selling substrate work that agents commoditise. They ship a tighter close, not authorship of the proprietary-data asset.

The finance-tech vendors will sell platform licences and a long implementation. Your proprietary-data moat has to sit on firm-owned infrastructure or it becomes a vendor asset at renegotiation. The CFO who builds capital intelligence on a vendor stack builds it for the vendor.

The FP&A and planning specialists will sell a point solution for one layer: scenario planning, close-acceleration, reporting. None can run the capital-allocation-plus-narrative conversation the seat requires.

The AI-native finance-agent startups will sell substrate tooling where depreciation is steepest. Buying them does not address who authors the capital-allocation decision. A cleaner close, the same authorship gap.

The market offers four things you have already said no to. GRAIL is the fifth option.

WHAT YOU GET

Senior advisory that speaks board-room, not deck-room

Working infrastructure your team uses the same week

Custom agents built on your project files, your book, your data

The start of a proprietary-data asset your FP&A owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried CFO-adjacent 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 finance leadership. Working infrastructure from week one.

We start with your finance 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 finance function. Build the first piece of infrastructure in the room: a live driver agent on your real renewal model, running against your real analyst CS, platform-usage, and procurement-benchmark data. Your leadership team leaves with a shared view of what to do next and the first piece of working infrastructure already in their hands.

WORKSHOP 02

Your finance infrastructure

A half-day build. Your Controller, Head of FP&A, Treasurer, and Head of IR each leave with five to ten pieces of working infrastructure tuned to their function. Subscription-economics agents. Continuous renewal scenario generation. Covenant-monitoring intelligence. Board-pack narrative drafting anchored in the proprietary-data asset.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The rollout plan to the wider finance organisation. The capital-allocation-thesis document for the next board meeting. The pricing-migration roadmap for the research and commercial teams. And the CEO and sponsor conversation with numbers they can check: ROIC alternatives, scenario ranges, proprietary-asset measurement.

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

For your next forecast review

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

Your CTO wants €2M for AI infrastructure. Is it in the IT budget as a productivity line, or on the capital page as one option among four, with proprietary-data systematisation on the same page? If the former, what would it take to move it before the next board cycle?

Your board voted on a 92% renewal plan. Your print is 84%. Can your team walk the audit committee through the account-level story in the same meeting, or does it take a post-mortem? If the latter, what data architecture is missing?

Your AI-native competitors are quoting on two or three pricing mechanics. You are on one. What is the twelve-month plan to migrate your custom-inquiry and advisory books off blended day rates, and who on your team is authoring it?

Your sponsor, your chair, and your CEO have each asked for a different AI-and-migration document in the last ninety days. Are they excerpting the same thesis anchored in the same proprietary-data asset, or are your teams drafting three documents that do not reconcile?

Of the junior analysts and Controllers you are cutting to agents this year, how many are the people who would have made your 2031 succession bench? What are you doing instead to build that bench?

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

▶ YOU ARE READING

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

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