

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in Your Fee Story

Your next €10M of firm capital. The cost-to-income ratio you sign. The thesis your founding partners can recite back. Three questions that look separate, but are actually one. A briefing for CFOs at mid-market European asset managers.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market
asset-management finance
leaders

DATE

April 2026

Monday morning, partner meeting Thursday

Every CFO at a mid-sized European asset manager has had the same Monday morning. Four things on the screen at 9:00. Q3 net flows printed at negative €320M against a flat-to-positive plan, the second consecutive quarter of outflows on the flagship strategy. Your lead PM on that strategy has a lift-out offer from a London competitor on the table, still open. Your CIO wants €2M next quarter to rebuild the research platform on a foundation-model stack, and his one-page says "investment technology infrastructure." Your gatekeeper consultant forwarded a systematic-peer DDQ packet on Sunday night with a single line: "worth a look before our call Thursday."

You have signed fifteen cost-to-income forecasts for this partner meeting. The pattern used to be simple. Hold the fee rate, hold the cost base, ride the market beta. Close the quarter clean. Keep the depositary happy. The business was understandable.

Something changed. Your gross fee rate is compressing five basis points a year and your cost base is not. Your AI capital is running through the trading-tech budget as a platform licence. Your AI-native systematic peers are running at one-fifth the PM count on comparable mandates. Your NAV strikes on time, and the fund-board oversight committee is asking about concentration risk on the flagship. The NAV is clean. The economics underneath it are not.

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 carrying 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 CIO 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 partner meeting.

Mid-market active just lost the equation that priced the industry for a decade

Every vanilla mandate your distribution team bids on now gets compared line by line against a passive equivalent at ten to fifteen basis points, a systematic alternative at eighteen to twenty-five, and the consultant's AI-screened peer set. Your gross fee rate holds at forty-five for now. Your cost base does not move. Across the next three years the mechanical difference is eight to twelve points of cost-to-income ratio, and it is already showing up in the partner-meeting deck. The LP consultants running these comparisons are not hostile. They are armed. FundLP, CENTRL, DD360, and CFA Institute's Augmented LP framework have put AI-driven fund-screening tools on every gatekeeper desk in Europe, and the fund that cannot demonstrate systematic, repeatable process loses before the meeting opens.

Your competitors are not waiting. The AI-native systematic houses are running a senior PM per two to four billion AUM plus agent substrate, versus the traditional house running one senior per five hundred million to a billion plus two or three analysts. On comparable mandates they quote a twenty-basis-point management fee and a cost-to-income ratio under fifty percent. The consolidators at mid-market scale are buying boutique systematic capability as tuck-ins rather than trying to rebuild themselves. A smaller set of the more patient ones are building the firm-intelligence infrastructure that systematic peers cannot replicate in eighteen months, because the proprietary edge sits in integrated investment data, LP-relationship data, and cost-to-income data underneath it. They are the quiet compounders.

Your gross fee rate is compressing. Your cost base is not. The answer sits in firm-capital discipline the finance seat has not yet claimed, not in operations.

The sector is repricing underneath everyone at the same time. Passive AUM crossed active in 2024 and the flow gap is still widening. Mid-market active sits in the squeeze zone between ten-basis-point index products and twenty-basis-point systematic alternatives, and the firms that treat this as a distribution problem rather than a capital-and-cost problem are walking into a partner meeting they will not win.

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

Your capital allocation. Your CIO's €2M AI-research-platform request is sitting in the trading-tech budget, routed through the same mechanism you use to sign off on terminal subscriptions and market-data feeds. It is not a tooling request. It is the largest single firm-capital decision you will make this decade, and it deserves the same discipline you bring to a boutique tuck-in or a seed-capital commitment. The CFOs who ship this as a firm-capital line to the founding partners, with alternatives analysis and an ROIC hurdle, keep the authorship for the next ten years. The ones who sign the requisition hand the authorship to the CIO and the Head of Trading Tech by default.

Your cost-to-income forecast. The team built the annual plan on a 62% cost-to-income target. The Q3 print came in at 67%. Your Head of FP&A cannot reconstruct the strategy-level story because the fee-compression signal lives in platform-level gross-to-net data, the flow-churn signal lives in distribution CRM notes, and the PM-capacity signal lives in attribution analysis 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 monthly management accounts, not for continuous reconstruction of the cost-to-income story.

Your fee architecture. The founding partners voted on a fee strategy that assumed vanilla active would hold its rate against passive. It isn't. Your systematic and AI-native peers are running redesigned performance-fee structures, commitment-tiered management fees, and outcome-linked pricing on newer strategies. You are running one club: flat management fee on AUM. The fee architecture that ran clean from 2014 to 2022 is now the reason your gross fee rate is compressing faster than your cost base. Fee architecture lives on the firm-capital page, where the moat gets built, and not on the distribution team's roadmap.

Your fund-accounting and FP&A bench. The NAV reconciliations, standard variance decomposition, depositary reporting, and first-draft DDQ generation that used to train a junior fund accountant into a Senior Controller into a Group Financial Controller are now running on agents. The Q1 headcount win is real. The 2031 succession bench is empty. The firm that cuts the junior layer without rebuilding the entry pathway has a cheaper quarter this year and no internal CFO candidate in five.

Your economic thesis. Your founding partners have asked for the strategy-offsite preview. Your fund-board chair has asked for a concentration-risk note on the flagship. Your lead consultant has asked how the firm stays relevant against systematic peers. 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 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 passive and systematic bring three pricing mechanics to the table and the rest of the market is on one. Your finance function is stretched because the cost-to-income credibility, 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 mid-market asset-management CFOs have not named yet, and it is the one that matters most.

Capital allocation without cost-to-income credibility is guessing. Cost-to-income credibility without a narrative is scorekeeping. A narrative without a firm-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 firm-intelligence infrastructure that feeds all three from the same data, the same governance, the same discipline.

This is the program-trading decision of this decade for active managers. Every CFO in your cohort is making it right now, whether they see it or not. The firms that built systematic capability in the 1995-2005 window kept a capital-allocation advantage for the next twenty years. The ones who treated it as an IT decision watched a peer become a different kind of firm. 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 partner-meeting question by walking the same firm-intelligence data, in the same meeting, with the same numbers? Most mid-market asset managers cannot. The ones that can have already made the shift, and they are the ones their founding partners are already betting the next decade on.

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 founding partners replace when the next generation walks in with a thesis of their own.

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

Most firms 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 trading-tech budget

Sign the CIO's requisition. Platform line in the plan. Trading tech owns the roadmap in eighteen months.

Hold the fee rate, absorb AI costs

Gross fee rate compresses five basis points a year. Cost base holds. Cost-to-income drifts four points above plan.

Run cost-to-income on a rolling spreadsheet

Monthly accounts. Three scenarios take a week. Flow-negative quarters produce no strategy story until the post-mortem.

Cut the fund-accounting layer for Q1

Headcount down twenty percent. Empty succession bench in 2031. External CFO hire the year after.

THE WINNERS' PATH

Route AI capital through the firm-M&A hurdle

Alternatives against tuck-ins, seed, distribution, partner returns. Founding partners see one firm-capital line.

Migrate fee architecture to three mechanics

Redesigned performance fee, commitment-tiered management fee, outcome-linked on newer strategies. Fee capture reflects AI alpha.

Rebuild cost-to-income as a live driver graph

Ten scenarios in an afternoon. Flow-negative quarters produce the strategy-level story in the same partner meeting.

Redesign the junior pathway around judgment

New hires pair with agents on scenarios, capital, DDQ work. Bench rebuilds faster than it would have grown.

Default shows a cleaner Q1. Winners have friction for three to five quarters and compound from month fifteen. The difference is whether the CFO authors the firm-capital decision or administers it.

Capital Allocation. Forecast Credibility. Value-Creation Narrative.

Three questions every asset-management 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 of firm capital actually compound?

AI-native alpha infrastructure? Passive-adjacent product launch with seed? Systematic-investing platform? LP-intelligence architecture? Boutique tuck-in? Distribution investment? Partner returns? Today's default: the CIO asks, the trading-tech 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 cost-to-income print above plan, and what is the strategy-level story?

Your founding partners voted on 62%. Your firm is printing 67%. Your Head of FP&A cannot walk the fund-board oversight committee through the strategy-level story because the fee-compression signal sits in platform-level gross-to-net data, the flow-churn signal sits in distribution CRM, and the PM-capacity signal sits in attribution her 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 founding partners want the strategy-offsite preview. Your fund-board chair wants the concentration-risk note. Your lead consultant wants the how-we-stay-relevant answer. Three audiences, three documents, three teams, and the numbers underneath them do not match. The CFOs pulling ahead write the thesis once and refresh it continuously from the same data architecture that runs the cost-to-income forecast. The audiences read excerpts of the same document.

What the firms winning this play are actually doing

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

1

Capital: route the AI-research decision through the firm-M&A hurdle, not the trading-tech requisition

Your CIO's €2M request is a firm-capital decision. Treat it like one. Present the alternatives on one page: AI-native alpha build, passive-adjacent product with seed, systematic-platform investment, boutique tuck-in, distribution build, partner returns. ROIC-scored against the same five-year horizon you use for a seed commitment. 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 cost-to-income as a live driver graph, not a rolling spreadsheet

Your cost-to-income model should draw from live AUM, live net-flow data, live gross-to-net fee data, and live PM-capacity attribution as one graph rather than five spreadsheets stitched together. Three fee-compression scenarios compress from a week to an afternoon. Flow-negative quarters produce the strategy-level story in the same partner meeting, not in the post-mortem. The forecast stops being a monthly deliverable and becomes a live artefact the CFO invests in continuously.

3

Fees: migrate from one management-fee club to three

Redesigned performance-fee architecture on the strategies that earn it. Commitment-tiered management fees for large institutional mandates with lock-ups. Outcome-linked or absolute-return-linked pricing on newer strategies where AI-enabled alpha is measurable. The flat management-fee rate holds the base. The three new mechanics convert systematic-investing capability into fee capture on the mandates where AI-native peers are already quoting. Finance authors the migration. Investment team and distribution build 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 firm capital gets starved to protect it? What gets fed? Refresh the thesis continuously off the same driver graph that runs the cost-to-income forecast. Partner-meeting deck excerpts the thesis. Fund-board oversight brief excerpts the thesis. Consultant update excerpts the thesis. Three audiences, one artefact, four registers. The CFO authors the narrative. The founding partners recite it back.

The only firm that joins capital authorship to working infrastructure

The asset-management CFO has four places to go. None can do what we do.

The big asset-management practices will sell a twelve-to-twenty-four-month cost-out programme staffed by a partner and analysts billing hours. Their audit arm constrains what they can advise on fund-board governance. Their billable-day model requires selling substrate that agents commoditise. They ship a tighter cost base, not authorship of the firm-capital decision.

The asset-management technology platforms will sell seat licences and a long implementation on their stack. Your firm-intelligence has to sit on firm-owned infrastructure or it becomes a platform asset at renegotiation. The CFO who builds firm-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, management reporting. None can run the capital-allocation-plus-narrative conversation the seat requires.

The systematic-investing and LP-tech vendors will sell substrate tooling where depreciation is steepest. Buying them does not address who authors the firm-capital decision. A cleaner DDQ response, 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 partner-meeting, not deck-room

Working infrastructure your team uses the same week

Custom agents built on your AUM book, your LP data, your cost base

The start of a firm-intelligence layer 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 cost-to-income driver agent on your real AUM book, running against your real fee-rate, flow, and gross-to-net 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 Group Financial Controller, Head of FP&A, Treasurer, and Head of IR each leave with five to ten pieces of working infrastructure tuned to their function. NAV and close agents. Continuous scenario generation on fee compression. ICAAP-and-regulatory-capital intelligence. DDQ and fund-board pack drafting. Infrastructure your team is using the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The rollout plan to the wider finance organisation. The firm-capital-thesis document for the next partner meeting. The fee-migration roadmap for the investment team and distribution. And the founding-partner and consultant conversation with numbers they can check: ROIC alternatives, scenario ranges, moat 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 partner meeting

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

Your CIO wants €2M for AI research infrastructure. Is it in the trading-tech budget as a platform line, or on the firm-capital page as one option among six? If the former, what would it take to move it before the next partner meeting?

Your founding partners voted on a 62% cost-to-income plan. Your print is 67%. Can your team walk the fund-board oversight committee through the strategy-level story in the same meeting, or does it take a post-mortem? If the latter, what data architecture is missing?

Your systematic and AI-native peers are quoting on three or four fee mechanics. You are on one. What is the twelve-month plan to migrate from one management-fee club to three, and who on your team is authoring it?

Your founding partners, your fund-board chair, and your lead consultant have each asked for a different AI-and-strategy document in the last ninety days. Are they excerpting the same thesis, or are your teams drafting three documents that do not reconcile?

Of the junior fund accountants and FP&A analysts 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 founding partners ask.

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