



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in Your LP Story

The fair-value mark you sign. The capital your GP deploys. The systematic edge your next fund is raised on. Three questions that look separate, but are actually one. A briefing for CFOs at mid-market European PE and VC funds.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market
PE and VC fund CFOs

DATE

April 2026

Thursday LPAC, DDQ due Monday, Managing Partner's €3M request

Every CFO at a mid-market European PE or VC fund has had the same Thursday morning. Four things on the screen at 8:30. The LPAC meets at two and three holdings are marked flat against a plan that called for a 1.2x step-up this quarter. The DDQ for the next fund raise came in Monday at fifty-two pages with a new operational section specifically on AI adoption and data governance. Your Managing Partner wants €3M next quarter for a portfolio-monitoring platform, and the one-page frames it as a technology line. Your newest LP on the committee emailed yesterday and asked for a Friday walkthrough on how the fund is using AI in diligence and value creation.

You have signed ten fund annuals. The pattern used to be simple. Close the quarter clean. Mark the book defensibly. Hold the LPAC. Keep the auditor happy. Raise the next fund on the track record. The business was understandable.

Something changed. Your distributions are running below the pacing model. Your auditor is asking deeper questions on fair-value judgment. Your LP operational due diligence team showed up last cycle with a questionnaire twice as long as the one before. Your portfolio CFOs are each making their own AI-infrastructure call with no shared playbook. Your Managing Partner is pitching the next fund on systematic edge, and the evidence underneath the claim sits in twelve different documents that have never been reconciled.

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 running since you took the seat are suddenly re-openable. The fund CFOs who treat them as engineering problems run the offensive play. The ones who treat them as permanent trade-offs sign requisitions in two directions and watch the next fund close at eighty percent of target.

This briefing covers what is happening in mid-market alternatives, what it means for the fund finance function you run, and the three questions your team should answer before the next Investment Committee.

Mid-market alternatives just lost the equation that raised the last three funds

Distributions are running at fourteen percent of NAV globally, the weakest since 2008-09. Hold clocks are shortening and the margin expansion that used to arrive at exit through multiple repricing is coming from inside the hold or not at all. Your LP's asset-allocation model is telling them to reweight away from funds whose J-curves are flatter than promised. Your track record still wins meetings. It no longer wins allocations on its own.

The DDQ has changed shape. The operational due diligence section now asks specific questions about sourcing systematisation, diligence process, portfolio-monitoring infrastructure, and AI adoption at both the GP and the portfolio. LPs are running these questionnaires against the answers you gave three years ago and comparing. The ones whose answers have visibly moved are winning the allocations. Every sponsor in your segment is claiming systematic edge. The allocator who runs eight of these questionnaires a quarter has learned to tell the difference between a narrative claim and an audit trail.

Your track record is still the price of entry. Your process evidence is now the price of allocation. Most funds cannot show the second without fabricating it.

The sector is repricing underneath everyone at the same time. The funds that put systematic infrastructure at the GP level in the 2015-2018 window ran their 2020-2023 raises on the asset they built. The ones that stayed on spreadsheets and personal memory spent the last cycle writing narrative into the deck and hoping the LPs did not audit it. In 2026 the next shift is compounding faster because the AI layer is ready, and the funds that move first are already visible in LP materials. The 2026-2028 fundraising cycle separates funds that can show process evidence at DDQ depth from funds that still cannot. The Managing Partners pulling ahead already know this. The CFOs who do not own the evidence layer hear about it in the LPAC.

Your marks, your fees, and your portfolio orchestration are feeling it every day

Your capital allocation at the GP. The €3M portfolio-monitoring-platform request sitting on your desk is not a technology line. It is the largest single capital-allocation decision the GP will make this fund cycle. The finance seats that treat it that way, with alternatives analysis against a GP-level AI build, a continuation-fund technology budget, and a reserve on management fees for the next raise, keep the authorship of the decision. The ones that sign the vendor contract hand the authorship to the portfolio-ops partner and the fund administrator by default, and the fund's systematic-edge story sits on someone else's infrastructure at the next DDQ.

Your fair-value marks. Three holdings are sitting flat at the LPAC. Your audit partner is asking how the scenario spread was built. Your fund administrator's template produced one number. Your own model produced a second. The LPAC's newest committee member is asking for a third using different comp-set assumptions. The mark is exposed because the data architecture underneath it was built for quarterly reconciliation. Nothing about it was designed to reconstruct the valuation story on demand, and the audit partner now expects that it can.

Your LP reporting and DDQ cycle. The quarterly pack consumed three-and-a-half weeks of fund-finance capacity last cycle. The DDQ for the next raise will consume six. The ILPA template stays stable and the LPs ask deeper questions underneath it every year. Your team cannot carry the LP reporting load and the DDQ load and the next-fund evidence layer with the same headcount. Either the close slips, the DDQ gets thin, or the headcount rises and the expense ratio the LPAC reviews shifts against you.

Your portfolio-CFO orchestration. Twenty-plus portfolio CFOs are each making their own AI-infrastructure call this year. Each one builds a different stack. Each one generates a different KPI format. Each one answers your quarterly data request differently. The fund that publishes a common playbook with shared infrastructure gets the portfolio graph the LPs want at DDQ. The fund that does not orchestrates twenty custom projects whose outputs never compound.

Your fund-finance bench. The NAV-close work, reconciliation, standard variance commentary, capital-account statement drafting, and first-pass DDQ response that used to train a junior fund accountant into a senior one into a Head of Fund Finance is now running on agents. The Q1 headcount win is real. The 2031 succession bench is empty. The finance function that cuts the junior layer without rebuilding the entry pathway has a cheaper quarter this year and no internal candidate in five.

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

Your sector is repricing from above because LPs are reading process evidence instead of track record. Your fund finance function is stretched because the fair-value marks, the GP-level capital allocation, the LP reporting cycle, and the portfolio-CFO orchestration are each being run 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 fund CFOs have not named yet, and it is the one that matters most.

Capital allocation without fair-value credibility is guessing. Fair-value credibility without a value-creation narrative is scorekeeping. A narrative without a data architecture underneath it is marketing. The fund-finance seat that holds all three holds one integrated asset, not three separate programmes. Every GP-level capital question, every fair-value defensibility gap, every DDQ weakness, every portfolio-orchestration gap is a symptom of the same missing layer: the capital-intelligence infrastructure that feeds all three from the same data, the same governance, the same discipline, across the fund and the portfolio.

This is the infrastructure decision of this fund cycle. Every fund CFO in your cohort is making it right now, whether they see it or not. The funds that put systematic GP-level infrastructure in place in the mid-1990s ran the 2004-2012 raises on the asset they built. The ones that stayed on spreadsheets spent the next decade playing catch-up. 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 a full fund cycle.

You can tell which camp a fund is in from a single question. Can your team walk an LP operational due diligence team through the fair-value process, the portfolio-monitoring graph, and the next-fund evidence layer, in one meeting, from the same data architecture? Most mid-market funds cannot. The ones that can are the ones winning allocations in the current cycle.

You have twelve to eighteen months. The 2026-2028 fundraising cohort is a short window. The CFOs who own this decision personally run the offensive play. The ones who do not become the CFO the Managing Partner replaces when the successor walks in with a thesis and a working prototype.

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

Every fund 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

Sign the portfolio-monitoring platform contract

Vendor stack. Technology line in the GP expense budget. Portfolio-ops partner owns the roadmap inside two quarters.

Mark the book from the fund-admin template

Quarterly NAV. One number per holding. LPAC debates the mark only when it moves ten percent.

Write the DDQ in fundraising-year Q3

Ninety-day scramble. Process answers drafted fresh. Evidence pulled from memory and selective documents.

Cut the fund-finance junior layer for the expense ratio

Headcount down. Tighter ratio this year. Empty succession bench in 2031. External hire that year.

THE WINNERS' PATH

Route GP AI capital through the fund hurdle

Alternatives on one page against the next-fund build, working capital, reserves. LPAC sees one capital line with options.

Rebuild marks as a live scenario process

Three scenarios per holding in an afternoon. LPAC sees the spread every quarter on every material position.

Build the DDQ evidence layer continuously

Refreshed from the same data architecture. DDQ responses become extracts, not projects. LP notices the depth.

Redesign the junior pathway around judgment

New hires pair with agents on fair-value, DDQ, and portfolio analysis. Bench rebuilds faster than it would have grown.

The default path shows a cleaner expense ratio in 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.

Capital Allocation. Forecast Credibility. Value-Creation Narrative.

Three questions every mid-market fund 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 GP capital actually compound, at the fund and across the portfolio?

Portfolio-monitoring platform? Fund-level AI infrastructure? Next-fund build? Reserves for follow-ons? The shared AI playbook that twenty portfolio CFOs will actually use? Today's default: the portfolio-ops partner asks, the GP budget funds, the vendor contract gets signed, the capital-allocation authority slides in two directions at once. The question is not which tool. The question is whose decision this is, and at which altitude.

02 · FORECAST CREDIBILITY

Why did the mark hold flat, and what is the portfolio-level story?

Your plan called for a 1.2x step-up this quarter. Three holdings held flat. Your audit partner asked how the scenario spread was built. Your fund administrator's template produced one number. The LPAC's newest committee member wants a third using different comp-set assumptions. Three valuation scenarios used to take a week. They now run in an afternoon. The bar has risen on what defensible and audit-ready means.

03 · VALUE-CREATION NARRATIVE

What does this fund do that compounds faster than peers?

Your next raise is pitched on systematic edge. Your DDQ lands Monday at fifty-plus pages. Your LP's operational team asks for process evidence at depth, across sourcing, diligence, portfolio operations, and AI orchestration. The funds pulling ahead build the evidence layer continuously from the same data architecture that marks the book. The DDQ becomes an extract. The LP meeting becomes a walkthrough. The next fund closes faster.

What the funds winning this play are actually doing

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

1

Marks: rebuild the quarterly fair-value close as a live scenario process

Your NAV close should draw from live portfolio operating data, live comp-set repricing, and live macro scenarios as one graph rather than four reconciliation templates stitched together the week before the LPAC. Three valuation scenarios per material holding compress from a week to an afternoon. The LPAC pack carries the scenario spread. The audit partner sees the reasoning. The fund-admin template becomes input, not output.

2

J-curve: rebuild the forecast as a live capital-call-and-distribution graph

Your J-curve projection, capital-call pacing, and distribution timing should refresh continuously from portfolio operating signals and the fund's own deployment model. Ten fund-cycle scenarios in an afternoon, not three per year. LPs see the expected-versus-actual pacing in every quarterly pack. The LPAC conversation shifts from explaining the gap to debating the forward choice.

3

DDQ: build the evidence layer continuously, not in fundraising-year Q3

Process, sourcing, diligence, value-creation, portfolio-monitoring, and AI-adoption evidence all drawn from the same architecture that runs the marks and the J-curve. DDQ responses become extracts. Operational due diligence teams are walked through living dashboards, not static decks. The next-fund close accelerates because the LP can verify at depth.

4

Portfolio: publish the AI-adoption playbook your portfolio CFOs run off

Common framework for portfolio-level AI infrastructure. Shared agents the fund CFO builds once. Shared data format where the legal structure allows. Each portfolio CFO adapts, not invents. The portfolio-monitoring graph compounds across twenty companies instead of fragmenting into twenty custom projects. Value creation accelerates inside the hold. The LP sees one systematic picture at DDQ.

The only firm that joins fund-level authorship to portfolio orchestration

The fund CFO has four places to go. None can do what we do.

The fund-administration platforms will sell a subscription and technology-enabled reporting. Their business model is per-fund operational services. The moat has to sit on firm-owned infrastructure or it becomes a vendor asset at renegotiation. The fund that builds systematic-edge evidence on an administrator's stack builds it for the administrator.

The portfolio-monitoring platforms will sell a dashboard and an implementation. They solve one layer cleanly and do not orchestrate capital at the GP or compound across fundraising cycles. The data graph lives in the vendor, not the fund.

The PE-focused strategy practices will sell a diagnostic engagement staffed by partners and analysts. They will write the LP narrative. They cannot put working fair-value agents or a DDQ evidence layer in your team's hands in three weeks.

The AI-native alternatives-finance startups will sell tooling where depreciation is steepest. A faster DDQ response or a slicker K-1 generator. Buying them does not build the fundraising-cycle asset underneath.

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

WHAT YOU GET

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

Working infrastructure your team uses the same week

Custom agents built on your fund, your book, your data

The start of an evidence layer your finance team 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 fund finance leadership. Working infrastructure from week one.

We start with the fund 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 the sector and inside your fund-finance function. Build the first piece of infrastructure in the room: a live fair-value scenario agent on one real holding, running against live comp-set and operating data. Your finance 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 fund infrastructure

A half-day build. Your Head of Fund Finance, your IR lead, and your Portfolio Operations partner each leave with five to ten pieces of working infrastructure. Fair-value scenario agents. J-curve projection modelling. DDQ evidence drafting. Portfolio KPI orchestration. Infrastructure your team uses the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The rollout plan to the wider fund team. The capital-allocation thesis document for the next Investment Committee. The AI-adoption playbook for your portfolio CFOs. And the LP conversation with evidence they can audit: fair-value reasoning, portfolio-level signals, DDQ extracts on demand.

Three weeks. No IT project. No Investment Committee approval needed to start. The infrastructure your fund finance team is already running on when we finish.

For your next Investment Committee

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

Your Managing Partner wants €3M for a portfolio-monitoring platform. Is it a technology line in the GP budget, or on the capital page as one option among four? If the former, what would it take to move it before the next IC?

Three holdings held flat at the last LPAC mark. Can your team walk the audit partner through three scenarios per holding in the same meeting, or does it take two weeks? If the latter, what data architecture is missing?

Your DDQ is going out in fundraising-year Q3. Is the evidence layer underneath it a continuously refreshed artefact, or a ninety-day scramble? Which of the two does your LP operational team already know how to spot?

Twenty-plus portfolio CFOs are each making AI-infrastructure calls this year. Do they run off a common playbook and shared infrastructure, or twenty custom projects? Who in your fund is authoring the shared layer?

Of the junior fund accountants and analysts you are rotating to agents this year, how many are the people who would have made your 2031 fund-finance bench? What are you doing instead to build that bench?

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