

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Margin Your Firm Sells

Your next €10M of capital. The project-margin 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 technology and digital firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market
technology and digital finance
leaders

DATE

April 2026

Thursday morning, board pack due next Thursday

Every CFO at a mid-sized European technology or digital services firm has had the same Thursday morning. Four things on the screen at 9:00. You lost the CTO pitch Friday to a twelve-person AI-native shop with a thirty-day proof-of-concept priced at a fraction of your six-month engagement. Your best solution architect resigned Monday for equity at an AI startup, and two more are in the same rolodex. Q3 utilisation printed at 52% against a 58% plan, with three fixed-price projects that ran hot. Your sponsor's partner forwarded a Bloomberg piece yesterday about an AI-native competitor that automated forty percent of delivery and cut headcount twenty percent, with one line: "curious to hear your thinking before Monday."

You have signed fifteen forecasts for this board. The pattern used to be simple. Hold the blended rate, hold utilisation at 58%, grow into the pipeline. Close the quarter clean. Keep the auditor comfortable on POC revenue recognition. The business was understandable.

Something changed. Your utilisation is drifting and your team cannot reconstruct the project-by-project story. Your AI capital is running through the CTO's tooling budget as a productivity line. Your AI-native competitors are pricing on outcomes, your rate card is priced per hour, and procurement teams are putting both quotes side by side in the CIO's inbox. Your Controllers are closing faster than ever, and the audit committee is asking about traceability on the agentic tooling your delivery team has been using since February. The close is clean. The margin story underneath it is 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 navigating since you took the seat are suddenly re-openable. The CFOs who treat them as pricing-mechanic 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 technology services just lost the pricing mechanic that funded the sector for a decade

Every AI coding assistant your delivery team rolls out raises engineer output by forty to sixty percent. Your blended rate on those engagements stays flat, because your pricing mechanic was built for a world where billable hours were what you sold. The AI-native entrants in your segment are not making this mistake. They ship the same capability priced per outcome, per deliverable, per resolved incident. Procurement teams at your best accounts are now running comparison quotes where your per-hour number sits next to their per-outcome number. You are running one pricing mechanic. They are running three. Across your installed base the mechanical difference is six to eight points of growth a year, and SEB's research already shows the print: 2.1% for time-based firms against 8.7% for value-based ones.

Your competitors are not waiting. Netcompany runs eighty percent fixed-price on eighty percent own-IP revenue, and it is outperforming the sector on profitability by a factor the rest cannot close on productivity alone. Sigma IT, Nexer, Knowit, and Bouvet are building productised services on top of their project books. The firms moving first are buying AI-native entrants as tuck-ins rather than trying to rebuild themselves. Some are migrating pricing on their AI-augmented engagements first and the rest of the book second. A smaller set of patient ones are building the delivery-intelligence moat that AI-native entrants cannot replicate in eighteen months.

Your engineer output is rising. Your revenue per engagement is not. That is not a delivery problem. It is a pricing-mechanic decision the finance seat has not yet claimed.

The sector is repricing underneath everyone at the same time. CIOs are approving thirty-day AI-assisted proofs-of-concept against what used to be six-month engagements. The credibility paradox compounds the pressure: a technology firm visibly behind on AI in a pitch is reputational damage, not just a lost deal. The firms that treat this as a sales-pitch problem rather than a pricing-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-delivery-infrastructure request is sitting in the tooling budget, routed through the same mechanism you use to sign off on laptops and Copilot licences. The request looks operational. The decision is the largest single capital-allocation call 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 and the CTO by default.

Your forecast. The team built the annual plan on a 58% utilisation and 12% project-margin assumption. The Q3 print came in at 52% utilisation and three fixed-price engagements ran hot. Your Head of FP&A cannot reconstruct the project-level story because the scope-creep signal lives in Jira, the attrition signal lives in Slack DMs, and the AI-automation pressure on pricing lives in competitor POC case studies her team has never seen. The plan was built on historical utilisation, and the data architecture underneath it was designed for monthly reconciliation, not for continuous reconstruction of the margin story when AI compresses delivery hours on half the book.

Your rate-card architecture. The board voted on a pricing strategy that assumed AI-assisted delivery would price like traditional delivery. It does not. Your AI-native competitors are on three pricing mechanics. You are on one. The rate card that ran clean from 2018 to 2024 is now the reason your utilisation is drifting on accounts where the comparison is clearest. Rate-card architecture reads as a sales question. In a market repricing underneath everyone, it becomes a capital-allocation question about where the moat gets built.

Your senior bench and junior pathway. Your best architects are being recruited by AI-native startups offering equity every week, and the Q3 attrition at senior level is running hot. The standard delivery work that used to train a junior consultant into a senior one (boilerplate code, unit testing, first-draft documentation, standard integrations) now runs 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 senior candidate in five.

Your economic thesis. Your board has asked for the strategy-day preview. Your sponsor's partner has asked for an AI-and-headcount note. Your chair has asked how the firm stays credible when AI-native competitors are pitching every week. 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 AI-native entrants bring three pricing mechanics to the table and the rest of the market is on one. Your finance function is stretched because the forecast 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 technology-services CFOs have not named yet, and it is the one that matters most.

Capital allocation without forecast credibility is guessing. Forecast credibility 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 delivery-intelligence infrastructure that feeds all three from the same data, the same governance, the same discipline.

This is the perpetual-to-SaaS pricing-mechanic migration of this decade. Every CFO in your cohort is making it right now, whether they see it or not. The ones who personally owned the 1996-2002 software-licence migration from perpetual to subscription kept capital-allocation authority for the next twenty years. The ones who signed the vendor licences because "that was a product decision" watched the board hire an external operator and the next CFO come from outside. 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 project-margin graph, in the same meeting, with the same data? Most mid-market technology services 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.

Most technology-services CFOs 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 CTO's tooling budget

Sign the Copilot and foundation-model requisition. Productivity line in the annual plan. CIO owns the roadmap within eighteen months.

Hold the rate card, ship AI delivery free

Engineer output rises forty to sixty percent. Blended rate flat. Utilisation drifts six points below plan. Project margin compresses.

Run the forecast on a rolling spreadsheet

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

Cut the junior layer for the Q1 win

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

THE WINNERS' PATH

Route AI capital through the M&A hurdle

Alternatives on one page against productised-IP, tuck-ins, senior-bench enablement, buyback. Board sees one capital line and one choice.

Migrate pricing to three mechanics

Fixed-price on productised services. Outcome-share on strategic retainers. Subscription on platform IP. AI efficiency converts to margin.

Rebuild the forecast as a live driver graph

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

Redesign the junior pathway around judgment

New hires pair with agents on scenarios, scope reading, client 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 pricing-mechanic migration or administering the tooling budget.

THREE QUESTIONS

Capital Allocation. Forecast Credibility. Value-Creation Narrative.

Three questions every technology-services 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?

AI-native delivery infrastructure?
Productised-IP investment?
Tuck-in of an AI-native shop? Senior-bench AI-enablement? Working-capital financing for the fixed-price migration?
Today's default: the CTO signs the tooling requisition, the capital-allocation authority slides away from the seat, and AI becomes a productivity line instead of a pricing-mechanic move. The question is not which tool. The question is whose decision this is.

02 · FORECAST CREDIBILITY

Why did utilisation print below plan, and what is the project-level story?

Your board voted on 58% utilisation. Your firm is printing 52% with three fixed-price projects that ran hot. Your Head of FP&A cannot walk the audit committee through the project-level story because the scope-creep signal lives in Jira, the attrition signal lives in Slack, and the AI-pricing-pressure signal lives in competitor POCs 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 AI-and-headcount note. Your chair wants the how-we-stay-credible answer when AI-native competitors are pitching every week. 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 architecture that runs the forecast. The three audiences read excerpts of the same document.

What the firms winning this play are actually doing

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

1 **Capital: route the AI-infrastructure decision through the M&A hurdle, not the CTO tooling requisition**

Your CTO's €2M request is a capital-allocation decision. Treat it like one. Present the alternatives on one page: AI-native delivery infrastructure, productised-IP investment, tuck-in of an AI-native shop, senior-bench AI-enablement, share 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 forecast as a live project-margin driver graph, not a rolling spreadsheet**

Your project-margin model should draw from live time-tracking, live scope-creep signal, live attrition data, and live pipeline-conversion benchmarks, as one graph rather than five spreadsheets stitched together. Three scenarios compress from a week to an afternoon. Below-plan quarters produce the project-level story in the same board meeting, not in the post-mortem. The forecast stops being a deliverable on a cadence and becomes a live artefact finance invests in continuously.

3 **Pricing: migrate from one rate-card mechanic to three**

Fixed-price on productised services where scope is known. Outcome-share or gain-share on strategic retainers where AI touches measurable client value. Subscription on platform IP that the firm owns. The per-hour rate card holds the commodity tail. The three new mechanics convert the forty to sixty percent engineer-output uplift into margin expansion on accounts where AI-native competitors are already quoting. Finance authors the migration plan. Delivery 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 when AI-native competitors pitch every week? What capital gets starved to protect it? What gets fed? Refresh the thesis 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 CFO authors the narrative. The board recites it back.

The only firm that joins pricing authorship to working delivery

The technology-services 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 that agents commoditise. They ship a tighter close, not authorship of the pricing-mechanic migration.

The PSA and finance-tech vendors will sell platform licences and a long implementation. They track utilisation and invoice cycles. None of them reprises your book. Your moat has to sit on firm-owned infrastructure or it becomes a vendor asset at renegotiation.

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-pricing-migration 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 book, your data, your IP

The start of a delivery-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 project-margin agent on your real book, running against your real time-tracking, pipeline, and client-health 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. POC-recognition agents. Continuous project-margin scenarios. Covenant-monitoring intelligence. Board-pack narrative 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 capital-allocation-thesis document for the next board meeting. The pricing-migration roadmap for the delivery team. And the CEO and sponsor 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 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-delivery infrastructure. Is it in the tooling budget as a productivity line, or on the capital page as one option among four? If the former, what would it take to move it before the next board cycle?

Your board voted on a 58% utilisation and 12% project-margin plan. Your print is 52% and three fixed-price projects ran hot. Can your team walk the audit committee through the project-level story in the same meeting, or does it take a post-mortem?

Your AI-native competitors are quoting on two or three pricing mechanics. You are on one. What is the twelve-month plan to migrate from one rate-card mechanic to three, and who on your team is authoring it?

Your sponsor, your chair, and your CEO 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 consultants and analysts you are cutting to agents this year, how many would have made your 2031 senior 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