



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Platform You Deliver From

Your delivery platform is the moat your firm has not yet named. Your R&D budget is two instruments pretending to be one. Your scorecard measures utilisation and developer productivity, and neither is the 2026 question. A briefing for CTOs, CPOs, Heads of Platform, and Heads of Delivery at mid-market European technology and digital services firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market technology
and digital services leaders

DATE

April 2026

Tuesday 10:00, delivery review

Every CTO at a mid-sized European technology-services firm has had the same Tuesday. Utilisation at sixty-five percent. Copilot adoption at seventy-two percent. Sprint velocity up sixteen percent on the dashboard. You lost a two-million-euro fixed-price pitch on Friday to an eleven-person shop that ran an AI-augmented delivery walkthrough in the pitch room. Your lead architect on the bank account mentioned a recruiter call from a twelve-person AI-native over coffee Thursday. Your CEO forwarded the SEB research this morning on value-based versus time-based growth. Your biggest client's CIO asked last week whether they could see your actual internal AI-augmented delivery workflow. You said yes. You are not sure what you are going to show them.

You have run this platform for ten years. The pattern used to be simple. Hire strong engineers. Hold utilisation. Deliver projects clean. Meet the board on billable throughput. The business was understandable.

Something changed. What used to take fifty developers on a delivery team now runs on five when the platform underneath them is designed for agents. That compression arrived faster than the discipline to absorb it. Ninety-five percent of corporate AI deployments produce no measurable profit impact, and the gap between what the tools can do and what your firm's system absorbs is where your seat now lives or loses. An eleven-person AI-native is pitching thirty-day POCs against your six-month engagements. Your senior architects carry delivery patterns that have never been written down, and two recruiters called last month.

You are not running one R&D function. You are running two, and only one of them is on your scorecard. One funds what you already deliver. The other funds what has to exist by 2028: the proprietary delivery platform you have not built, the productised services your client work has generated that you have never captured, the reusable-asset library that compounds across engagements, the integration depth that turns a six-month engagement into a thirty-day POC you can actually ship. 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 board cycle.

Technology services just lost the equation that priced the hour

The AI-natives in your category are not running pilots. They are pitching. An eleven-person shop that launched in 2024 just won a two-million-euro fixed-price against your seventy-person proposal. Their cost base is forty to sixty percent leaner. They demoed an AI-augmented delivery walkthrough in the pitch room that your firm has no equivalent to. The client procurement conversation has changed underneath everyone at the same time: CIOs are approving thirty-day AI-assisted POCs that deliver what used to require a six-month engagement. Your sales cycle was a moat. It is a liability now.

The credibility paradox is sharpest in this sector. Every client CTO walks your delivery floor expecting to see the firm that advises them on transformation running its own AI-augmented stack internally. What most see is a delivery environment unchanged from three years ago. In engineering or professional services that reads as operational inefficiency. In technology services it reads as reputational damage. The firm that shows up to the pitch with a demoable internal workflow wins the conversation before price enters it. The one that talks about AI theoretically loses to the one that ran the walkthrough.

Nordic companies spend forty to fifty percent of their AI budget on off-the-shelf productivity tools versus eight to eleven percent globally. Over-invested in generic tools. Under-invested in the domain-specific capability that compounds.

Underneath the pitch-room conversation is the pricing-mechanic migration. SEB research reads plainly: time-based firms grew 2.1 percent last year, value-based and productised firms grew 8.7 percent. A factor of four. Netcompany anchors the destination at eighty percent own-IP and eighty percent fixed-price, outperforming time-based peers on profitability with the same client base. The firms that repriced around productised services and outcome-priced engagements captured the AI efficiency as margin. The firms that kept billing hours watched clients demand the efficiency back as a discount. The delta compounds every quarter. The gap between a time-based firm and a value-based firm on any peer comparison a board chair reads this year is wider than last year and wider still next year.

Your delivery platform is feeling it every day

Your build engine. Copilot adoption is at seventy percent. Cycle time has not moved. Velocity is up double-digit on the dashboard and your reviewers are pushing ten-thousand-line PRs that are barely reviewed. The foundational layer rots underneath the green metrics. Twelve months in, your team ships faster-but-shallow and cannot execute a platform redesign when one arrives. The review discipline that worked in 2022 is now training the wrong practice on agent-generated volume.

Your delivery platform and data. Spotify's best developers have not written a line of code since December because Spotify spent four years building the platform foundation that lets agents ship production-grade. You cannot replicate four years in eighteen months. You can buy the build-side discipline from the tool layer and redirect the compounding capital into the layer the tool vendors cannot sell you: the firm-owned AI-augmented delivery tooling your sales team demos in every pitch, the reusable-asset library with editorial discipline, the integration-depth graph for the top ten client systems. The tool layer commoditises. The platform you own compounds.

Your product thesis and productised services. Netcompany hit eighty percent of revenue from own IP. You are at three. Your last four client engagements each solved a variant of the same problem, and the resulting code lives in four separate repos that nobody has touched since delivery closed. Your rate card assumes AI features will price like legacy features. They do not. The joint Pricing Architecture Document with your CFO and CRO needs to be written before the first productised-service engagement closes, not after.

Your senior bench. You stopped hiring juniors in 2024. The routine code-writing that trained juniors into mid-levels is now running on agents. No routine work, no training ground, no senior bench in 2030, no internal CTO or Head of Practice candidate the year after. OpenAI and Anthropic both resumed junior hiring in late 2025 specifically to pair juniors with seniors on judgment-intensive work. The firm that cuts juniors for this year's utilisation number has a cheaper function in Q4 and an empty bench in five years, against AI-native shops that are paying your top five people in equity.

Your board-level authority. Thirty percent of European AI-strategy ownership sits with the CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO. You are not on the list. Not because you abandoned the mandate. You have not yet claimed it. Your CEO has a CAIO shortlist. Every month without a visible claim, the platform-level thesis drifts further from the seat that can defend it. Your Q1 boundary agreement with the CEO lands before the CAIO shortlist, not after.

The moat that matters in 2028 is not the developers on your bench

Your sector is repricing from below because AI-native shops are pitching thirty-day POCs against your six-month engagements and demoing AI-augmented delivery in the pitch room. Your function is compressing because five developers on the right platform now produce what fifty did, and the review bar that worked for fifty cannot hold the volume five produce with agents. Same force, two altitudes. But there is a third thing happening that most mid-market technology-services CTOs have not named yet.

The moat that matters in 2028 is not the developers on your bench. It is the platform they deliver from and the IP your client work has generated that you have not yet turned into an asset. The developer-productivity tools commoditise. The model layer commoditises. What does not commoditise is the proprietary delivery platform you own, the productised services the firm sells as repeatable offerings, the reusable-asset library every new engagement walks into, and the integration depth your delivery has into the top ten client systems. An eleven-person AI-native shop cannot replicate that because they do not have your client base, your delivery history, or your installed integration. You do. The window closes in roughly eighteen months, which is the time between now and when the AI-native has enough delivery footprint of its own to start building the same layer.

In 2007, a company ran a DVD-by-mail warehouse. The tech stack was designed for fulfilment. The CTO's scorecard was warehouse uptime and mail-cycle turnaround. Every metric was green. The CEO bet the company on rebuilding the entire backbone for streaming. The investment read like capex pain for two years. By 2013 they were shipping House of Cards on infrastructure a peer's CTO in 2007 could not have funded under his own board's scoring unit. The peer had a functional online-rental business by 2004 and could not walk the sixteen-percent-of-revenue late-fee line-item away from his board. The 2007 scorecard never turned red. It just stopped being the one the market was scoring.

You can tell which camp a firm is in from a single question. Is the R&D-and-platform budget you defended at the last board meeting a plateau budget, a compounding budget, or one budget pretending to be one when it is two? Most mid-market incumbents cannot answer. The ones that can have already made the shift.

You have twelve to eighteen months. The CTOs who author the platform decision in Q1 author the decade. The ones who wait inherit a narrowed seat.

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

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

THE DEFAULT PATH

Ship Copilot, track adoption

Seventy-percent adoption. Cycle time flat. Foundation rots under a green dashboard.

Route AI capital through the IT requisition line

CIO owns the platform within eighteen months. CAIO arrives in Q3. CTO narrows to uptime.

Keep the AI efficiency internal as margin

Clients notice. Efficiency gets demanded back as a discount. Margin compresses anyway.

Cut the junior layer for the Q1 headcount win

Cheaper function this quarter. Empty succession bench by 2031. External CTO hire at a premium.

THE WINNERS' PATH

Buy the tool layer, redesign the review bar

Review depth per line. ADRs mandatory on agent-generated work. Eval harness as first-class.

Route platform capital through the M&A hurdle

Two budgets, two hurdle rates, two kill criteria. Plateau delivers margin. Compounding delivers the 2028 platform.

Ship one productised service on protected P&L

Fixed price, named outcome, demoable workflow. Parallel to the billable book. No cannibalisation.

Pair juniors with seniors on compounding work

OpenAI and Anthropic resumed junior hiring in late 2025. Eighteen-month apprenticeship on judgment.

Same delivery team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the platform capital or signs the requisition.

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

Three questions every technology-services 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 engineering speed shipping client-ready fixed-price output, or demos that fall over?

Copilot adoption at seventy percent. Cycle time flat. Ten-thousand-line PRs barely reviewed. Your clients see what you run internally the first time they walk your delivery floor. Spotify spent four years on the platform that lets agents ship cleanly. You cannot replicate four years in eighteen months. Buy what the tool layer can carry. Own the discipline that makes it ship.

02 · PRODUCT DEFENSIBILITY

What does our delivery produce that an eleven-person AI-native shop cannot copy in a year?

Your individual engineers are recruitable. Your proprietary delivery platform, your productised services, your reusable-asset library, and your integration depth into the top ten client systems are not. The moat sits in the platform, not on the bench. Netcompany ran this to eighty percent of revenue from own IP. Every firm that did not is headcount arbitrage against the AI-native shop.

03 · R&D CAPITAL ALLOCATION

Is our R&D-and-platform budget one instrument or two?

Plateau capital runs the existing billable engine at lower unit cost. Compounding capital builds the platform and productised IP for 2028. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The CTO who walks into the board with one budget runs the programme every peer is running. The one who walks in with two, each defended separately, authors the decade.

What the firms winning this play are actually doing

Four concrete moves any technology-services 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-generated velocity without the foundation rotting**

Instrument review-time-per-line as a floor, not just cycle time. PRs above five hundred lines require two senior reviewers and thirty minutes of documented review. Architectural decision records mandatory on every agent-generated change above a complexity threshold: no ADR, no merge. Evaluation harness coverage as a first-class metric on every fixed-price engagement: golden set, rubric, automated grading, no harness, no ship. The dashboard is no longer sprint velocity. It is review depth against output volume and client-ready delivery rate on fixed-price work.

2**Delivery platform and data: build the moat underneath the tool layer**

Three programmes that stack over eighteen months. A firm-owned AI-augmented delivery tooling stack your sales team demos in every pitch (the credibility-paradox response). A reusable-asset library with editorial discipline — structured, searchable, agent-accessible — capturing what the last three years of client work has actually generated. An integration depth graph for your top ten client systems, maintained as structured metadata your delivery team walks into every new engagement with. The tool vendors rent seats. This layer compounds.

3**Product thesis: stand two productised services on protected P&L**

Do not migrate the whole book. Pick the three most-repeated client problems from the last twenty-four months. Package two into named, fixed-price, outcome-specified offerings with demoable internal workflows. Ship on a separate P&L parallel to the billable book. Write the cannibalisation trigger and kill criteria in Q1. The product-and-platform architecture supports both pricing models natively. The joint Pricing Architecture Document with your CFO and CRO is signed before the first productised engagement closes.

4**Senior bench: rebuild the junior pathway around senior and agent pairing**

Your 2024 decision to stop hiring juniors saves Q1. It empties the 2031 senior bench. Reverse it. Juniors in 2026 do not start on routine code; the agent layer absorbed that work. They start on exception-handling, eval-harness maintenance, and paired composite delivery with a named senior for their first eighteen months. The career curve compresses from five years to eighteen months because the senior's judgment is now legible in the agent-and-junior review loop. By 2029 you have a mid-tier bench the AI-native shop down the street cannot hire away because it did not exist anywhere else to hire from.

The only firm that joins platform authorship to working delivery

The technology-services CTO has four places to go, and none can do what we do.

The big tech-transformation practices ship a standardised framework scaled across a sector book. Their cost structure requires reusability. Your delivery platform is idiosyncratic. The output is governance theatre you cannot defend to the board.

The hyperscaler professional services arms price engagements as enablement for cloud consumption. Their business model requires platform lock-in. Every build-or-buy decision routes through their pricing for three years.

The developer-productivity tooling vendors have a quote on your desk. Cursor, Copilot enterprise, Claude Code. Their product is a seat-measured velocity intervention. Spending freed capacity on the proprietary-platform moat instead of expanded seats is impossible for them to recommend.

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

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

WHAT YOU GET

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

Working infrastructure your team uses the same week

Custom agents built on your stack, your clients, your methodology

The start of a platform-and-productisation moat your firm owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried CPO-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 platform leadership. Working infrastructure from week one.

We start with your R&D and delivery-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 delivery platform. Build the first piece of infrastructure in the room: a platform-and-productisation capital-allocation agent that takes your real R&D budget, splits it two ways, and runs the alternatives analysis against your named AI-native competitors. Your leadership team leaves with a shared view and the first working artefact already in hand.

WORKSHOP 02

Your delivery infrastructure

A half-day build. Your Head of Delivery, Heads of Practice, Head of Platform, and senior architects each leave with five to ten pieces of working infrastructure tuned to their function. Reusable-asset library scaffolding. Review-bar audit agent. Productised-service scoping tool. Integration-depth graph prototype. A demoable internal workflow your sales team can walk a client CTO through. Infrastructure your team is using the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The board document that splits the R&D-and-platform budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the CEO on who authors the AI thesis. The joint Pricing Architecture Document with your CFO and CRO. The eighteen-month platform and productised-services build plan.

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

For your next delivery-platform review

The questions where your R&D and delivery leadership cannot agree on an answer are the ones worth an hour on the agenda. Ask these honestly.

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

Name the AI-native shop that just won the last fixed-price pitch you lost. What was in their demo that yours did not have? If your sales team cannot name the top three, that is the first thing the briefing changes.

If your two most senior architects left for an AI-native shop tomorrow, how many months would it take to reconstruct the delivery patterns behind your five biggest engagements? Is that reasoning captured in the platform, or does it live in their heads?

Copilot adoption is at seventy percent. Cycle time has not moved. Where did the freed hours go, and is your review bar still the one you want the foundational layer of your delivery platform trained on?

Heidrick's February 2026 data: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO, zero percent CTO on the European AI-strategy-ownership list. Is your Q1 boundary agreement with the CEO 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 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

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